

HP Reliable Transaction Router

Getting Started

Order Number: AA-RLE1C-TE

June 2005

This document introduces HP Reliable Transaction Router and describes its concepts for the system manager, system administrator, and applications programmer.

Revision/Update Information: This manual supersedes *Reliable Transaction Router Getting Started*, Version 4.2.

Software Version: HP Reliable Transaction Router Version 5.0

Hewlett-Packard Company
Palo Alto, California

© Copyright 2003, 2005 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group. Java is a US trademark of Sun Microsystems, Inc.

This document was prepared using DECdocument, Version 3.3-1B.

Contents

Preface	vii
1 Introduction	
Reliable Transaction Router	1-1
RTR Continuous Computing Concepts	1-3
RTR Terminology	1-4
RTR Server Types	1-17
RTR Networking Capabilities	1-26
2 Architectural Concepts	
The Three-Tier Architecture	2-1
RTR Facilities Bridge the Gap	2-3
Broadcasts	2-4
Flexibility and Growth	2-4
Transaction Integrity	2-5
The Partitioned Data Model	2-6
Object-Oriented Programming	2-6
Objects	2-8
Messages	2-9
Class Relationships	2-9
Polymorphism	2-10
Object Implementation Benefits	2-10
Java Support	2-11
XA Support	2-11

3 Reliability Features

RTR Server Types	3-1
Failover and Recovery	3-2
Router Failover	3-2
Recovery Scenarios	3-3
Backend Recovery	3-3
Router Recovery	3-3
Frontend Recovery	3-4

4 RTR Interfaces

RTR Management	4-6
RTR Administrator	4-6
RTR Command Line Interface	4-10
Application Programming Interfaces	4-16
RTR Java Object-Oriented Interface	4-16
RTR C++ Object-Oriented Programming Interface	4-18
RTR C Programming Interface	4-21

5 The RTR Environment

The RTR System Management Environment	5-1
Monitoring RTR	5-3
Transaction Management	5-4
Partition Management	5-5
RTR Runtime Environment	5-6
What's Next?	5-7

Glossary

Index

Examples

2-1	Objects-Defined Sample	2-9
4-1	The user issues the following commands on the server application where RTR is running on the backend.	4-12
4-2	When the next command is issued, RTR waits for the message from the client, which does not appear until after the client sends it (Example 4-3).	4-13

4-3	Commands and system response at client.	4-14
4-4	The following lines arrive at the client from RTR after the user enters commands at the server.	4-15
4-5	Creating a Facility with the C++ API	4-20
4-6	Starting RTR with the C++ API.	4-21
4-7	Creating a Facility with the C++ API	4-21
4-8	Creating a Partition with the C++ API	4-21

Figures

1	RTR Reading Path	xi
1-1	Client Symbol	1-5
1-2	Server Symbol	1-6
1-3	Roles Symbols	1-7
1-4	Facility Symbol	1-8
1-5	Components in the RTR Environment	1-8
1-6	Two-Tier Client/Server Environment	1-11
1-7	Three-Tier Client/Server Environment	1-11
1-8	Browser Applet Configuration	1-12
1-9	RTR with Browser, Single Node, and Database	1-13
1-10	RTR Deployed on Two Nodes	1-13
1-11	RTR Deployed on Three Nodes	1-14
1-12	Standby Server Configuration	1-15
1-13	Transactional Shadowing Configuration	1-16
1-14	Standby Servers	1-19
1-15	Shadow Servers	1-20
1-16	Concurrent Servers	1-21
1-17	A Callout Server	1-22
1-18	Bank Partitioning Example	1-24
1-19	Standby with Partitioning	1-25
2-1	The Multitier Model	2-2
2-2	Partitioned Data Model	2-7
4-1	RTR Administrator	4-2
4-2	RTR Command Line Interface	4-3
4-3	RTR Manager.	4-7
4-4	RTR Explorer: View of All Facilities	4-8
4-5	RTR Explorer: View of One Facility	4-9

4-6	RTR Service Provider	4-17
5-1	RTR System Management Environment	5-4
5-2	RTR Runtime Environment	5-7

Tables

1	RTR Documents	viii
2-1	Functional and Object-Oriented Programming Compared . . .	2-8
4-1	RTR Development Interfaces and their Use	4-5

Preface

Purpose of this Document

The goal of this document is to assist an experienced system manager, system administrator, or application programmer to understand the Reliable Transaction Router (RTR) product.

Document Structure

This document contains the following chapters:

- Chapter 1, Introduction, provides information on RTR technology, basic RTR concepts, and RTR terminology.
- Chapter 2, Architectural Concepts, introduces the RTR three-tier model and explains the use of RTR functions and programming capabilities.
- Chapter 3, Reliability Features, highlights RTR server types and failover and recovery scenarios.
- Chapter 4, RTR Interfaces, introduces the management and programming interfaces of RTR.
- Chapter 5, The RTR Environment, describes the RTR system management and runtime environments, and provides explicit pointers to further reading in the RTR documentation set.

Related Documentation

Table 1 describes RTR documents and groups them by audience.

Table 1 RTR Documents

Document	Content
For all users:	
<i>HP Reliable Transaction Router Release Notes</i> ¹	Describes new features, corrections, restrictions, and known problems for RTR.
<i>HP Reliable Transaction Router Getting Started</i>	Provides an overview of RTR technology and solutions, and includes the glossary that defines all RTR terms.
<i>HP Reliable Transaction Router Software Product Description</i>	Describes product features.
For the system manager:	
<i>HP Reliable Transaction Router Installation Guide</i>	Describes how to install RTR on all supported platforms.
<i>HP Reliable Transaction Router System Manager's Manual</i>	Describes how to configure, manage, and monitor RTR.
For the application programmer:	
<i>HP Reliable Transaction Router Application Design Guide</i>	Describes how to design application programs for use with RTR, with both C++ and C interfaces.

¹Distributed on software kit.

(continued on next page)

Table 1 (Cont.) RTR Documents

Document	Content
<i>HP Reliable Transaction Router JRTR Getting Started</i> ²	Provides an overview of the object-oriented JRTR Toolkit including installation, configuration and Java programming concepts, with links to additional online documentation.
<i>HP Reliable Transaction Router C++ Foundation Classes</i>	Describes the object-oriented C++ interface that can be used to implement RTR object-oriented applications.
<i>HP Reliable Transaction Router C Application Programmer's Reference Manual</i>	Explains how to design and code RTR applications using the C programming language and the RTR C API. Contains full descriptions of the basic RTR API calls.

²In downloadable kit.

You can find additional information about RTR, including the Software Product Descriptions, on the RTR website found through <http://www.hp.com> links to middleware products or at <http://www.hp.com/go/rtr> .

Reader's Comments

HP welcomes your comments on this manual. Please send comments to either of the following addresses:

Internet	rtrdoc@hp.com
Postal Mail	Hewlett-Packard Company RTR Documentation Group, ZKO3-4/Y17 110 Spit Brook Rd. Nashua, NH 03062-2698

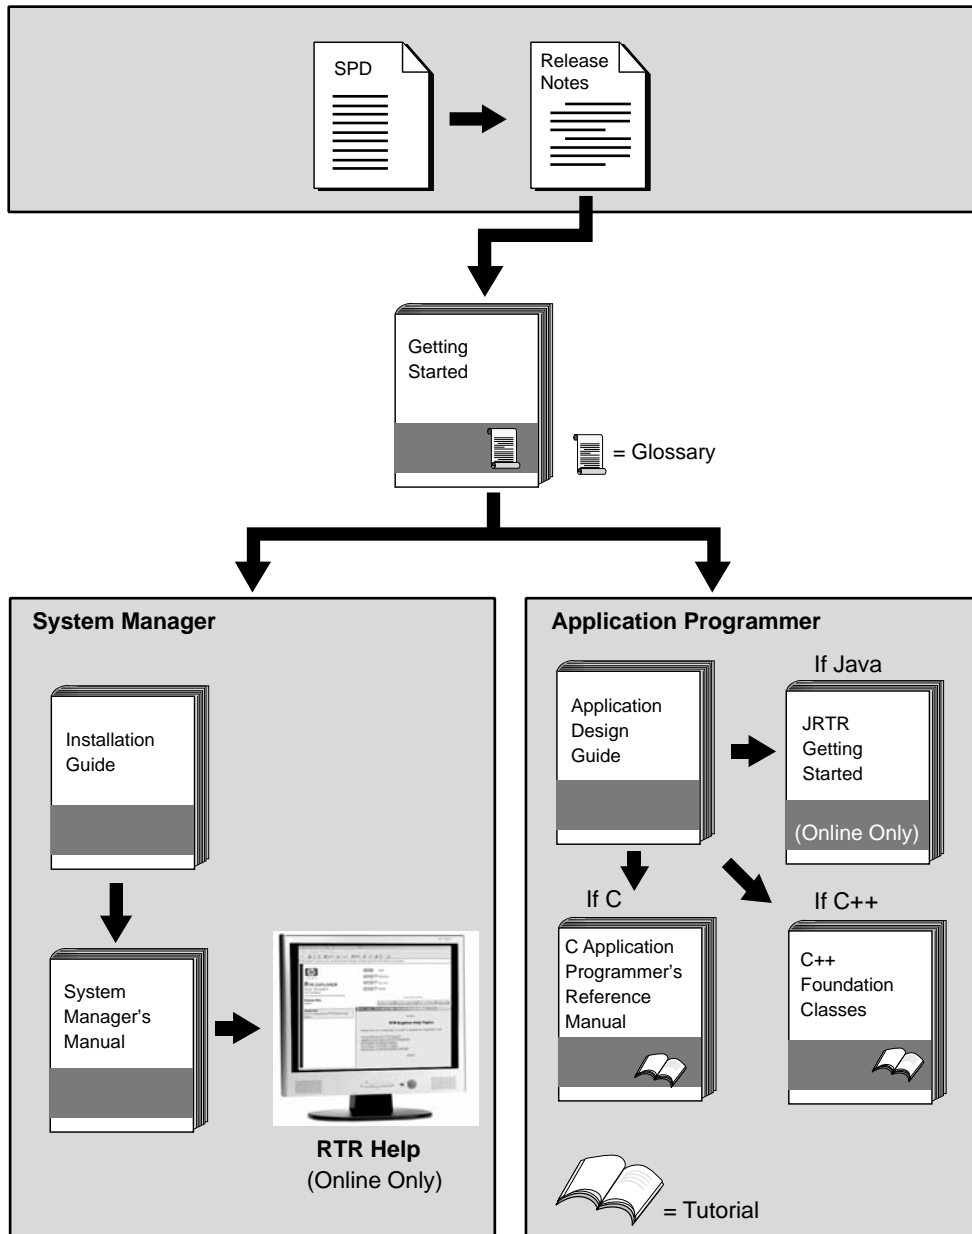
Conventions

This manual adopts the following conventions:

Convention	Description
New term	New terms are shown in bold when introduced and defined. All RTR terms are defined in the glossary at the end of this document.
User input	User input and programming examples are shown in a monospaced font. Boldface monospaced font indicates user input.
<i>Terms and titles</i>	Terms defined only in the glossary are shown in italics when presented for the first time. Italics are also used for titles of manuals and books, and for emphasis.
FE	RTR frontend
TR	RTR transaction router or router
BE	RTR backend
Refer to . . .	References another manual.
See . . .	References another part of this manual.

Reading Path

The reading path to follow when using the Reliable Transaction Router information set is shown in Figure 1.



VM-0818A-AI

1

Introduction

This document introduces RTR and describes RTR concepts. It is intended for the system manager or administrator and for the application programmer who is developing an application that works with Reliable Transaction Router (RTR).

Reliable Transaction Router

Reliable Transaction Router (RTR) is failure-tolerant transactional messaging middleware used to implement large, distributed applications with client/server technologies. RTR helps ensure business continuity across multivendor systems and helps maximize uptime.

Failure and fault tolerance

Failure tolerance is supplied by the RTR software that enables an application to continue even when failures such as node or site outages occur. Failover is automatic. **Fault tolerance** is supplied by systems with hardware that is built with redundant components to ensure that processing survives failure of an individual component. Depending on system requirements and hardware, an RTR configuration can be both failure and fault tolerant.

Reliable Transaction Router

Interoperability

You use the architecture of RTR to ensure high availability and transaction completion. RTR supports applications that run on different hardware and different operating systems. RTR applications can be designed to work with several database products including Oracle, Microsoft Access, Microsoft SQL Server, Sybase, and Informix. For specifics on operating systems, operating system versions, and supported hardware, refer to the *HP Reliable Transaction Router Software Product Description* for each supported operating system.

Networking

RTR can be deployed in a local or wide area network and can use either TCP/IP or DECnet for its underlying network transport.

Transaction processing

A transaction involves the exchange of something for something else, for example, the exchange of cash for something purchased such as a book or restaurant meal. Transactions are the fodder of the commercial world in which we live. Transaction processing is the use of computers to do the bookkeeping for the physical transactions in which people engage. RTR is the premier transaction processing software available for many computer systems, offering unique benefits to ensure the integrity and correctness of transactions under its control.

Corporations and institutions benefit when the data they need is current and kept rapidly up-to-date. This ensures that they have increased control of their business or institution, and can react more effectively to changes in the business or institutional environment.

For example, an inventory management system provides current inventory for inquiry and can automatically request inventory updates when needed. Or a production manager can make production decisions based on current status of manufacturing elements directly from the shop floor, and a credit card company can respond with accurate information when a card request for validation arrives. Or a shipping company can determine and report the location of a package in transit anywhere in the world, or a personnel system can enable an employee to update personal information online. All of these systems perform transaction processing tasks as designed by their developers.

With RTR, applications can be written to be deployed over a wide geography to take advantage of distributed resources, both computers and personnel. Implementing a transaction

processing system using RTR requires analysis, planning, and considered execution.

RTR Continuous Computing Concepts

RTR provides a continuous computing environment that is particularly valuable in financial transactions, for example in banking, stock trading, or passenger reservations systems. RTR satisfies many requirements of a continuous computing environment:

- Reliability
- Failure tolerance
- Data and transaction integrity
- Scalability
- Ease of building and maintaining applications
- Interoperability with multiple operating systems

RTR additionally provides the following capabilities, essential in the demanding transaction processing environment:

- Flexibility
- Parallel execution at the transaction level
- Potential for step-by-step growth
- Comprehensive monitoring tools
- Management station for single console system management
- WAN deployability

RTR also ensures that transactions have the ACID properties that have been established as crucial in a transaction processing environment. A transaction with the ACID properties is:

- Atomic
- Consistent
- Isolated

- Durable

For more details on transactional ACID properties, see the brief discussion later in this document in the section Transaction Integrity and refer to the *HP Reliable Transaction Router Application Design Guide*.

RTR Terminology

In addition to the terms previously defined, the following terms are either unique to RTR or redefined when used in the RTR context. If you have learned any of these terms in other contexts, take the time to assimilate their meaning in the RTR environment. These and other terms are also defined in the Glossary of this manual. The terms are described in the following order:

- Application
- Client, client application
- Server, server application
- Channel
- RTR configuration
- Roles
- Frontend
- Router
- Backend
- Facility
- Transaction
- Transaction controller
- Transactional messaging
- Nontransactional messaging
- Transaction ID
- Tier

- Standby server
- Transactional shadowing
- RTR journal
- Partition
- Key range
- XA

RTR Application

An RTR application is user-written software that executes within the confines of several distributed *processes*. The RTR application may perform user interface, business, and server logic tasks and is written in response to some business need. An RTR application can be written in one of the supported languages, C, C++, or Java and includes calls to RTR. RTR applications are composed of two kinds of actors, client applications and server applications. An application process is shown in diagrams as an oval, open for a client application (see Figure 1–1), filled for a server application (see Figure 1–2).

Client

A **client** is always a **client application**, one that initiates and demarcates a piece of work. In the context of RTR, a client must run on a node defined to have the frontend role. Clients typically deal with presentation services, handling forms input, screens, and so on. A client could connect to a browser running a browser *applet* or be a webserver acting as a gateway. In other contexts, a client can be a physical system, but in RTR and in this document, physical clients are called frontends or nodes. You can have more than one instance of a client on a node.

Figure 1–1 Client Symbol

VM-0819A-AI

RTR Terminology

Server

A **server** is always a *server application*, one that reacts to a client's units of work and carries them through to completion. This may involve updating persistent storage such as a database file, toggling a switch on a device, or performing another predefined task. In the context of RTR, a server must run on a node defined to have the backend role. In other contexts, a server can be a physical system, but in RTR and in this document, physical servers are called backends or nodes. You can have more than one instance of a server on a node. Servers interact with partitions and can have partition states such as primary, standby, or shadow.

Figure 1–2 Server Symbol



VM-0820A-AI

Channel

RTR expects client and server applications to identify themselves before they request RTR services. During the identification process, RTR provides a tag or handle that is used for subsequent interactions. This tag or handle is called an **RTR channel**. A channel is used by client and server applications to exchange units of work with the help of RTR. An application process can have one or more client or server channels. Channel management is handled transparently by the C++ and Java APIs.

RTR configuration


An RTR configuration consists of *nodes* that run RTR client and server applications. An RTR configuration can run on several operating systems including OpenVMS, Tru64 UNIX, and Windows among others (for the full set of supported operating systems, refer to the appropriate SPD). Nodes are connected by network *links*.

Roles

A node that runs client applications is called a **frontend** (FE), or is said to have the frontend role. A node that runs server applications is called a **backend** (BE). Additionally, the transaction **router** (TR) contains no application software but acts as a traffic cop between frontends and backends, routing transactions to the appropriate destinations. The router controls the distributed RTR nodes, and takes care of two-phase commit, failover and fallback.

The router also eliminates any need for frontends and backends to know about each other in advance. This relieves the application programmer from the need to be concerned about network configuration details. The router can reside on a node running as a frontend or a backend but is often run on a node where neither backends nor frontends are running. Figure 1–3 shows the symbol for each of the RTR roles.

Figure 1–3 Roles Symbols




VM-0821A-AI

Facility

The mapping between nodes and roles is done using a **facility**. An RTR facility is the user-defined name for a particular configuration whose definition provides the role-to-node map for a given application. The facility symbol (see Figure 1–4) illustrates its use in the RTR environment. Nodes can share several facilities. The role of a node is defined within the scope of a particular facility. Normally a facility is defined across all roles but facility definition depends on application design.

The router is the only role that knows about all three roles. A router can run on the same physical node as the frontend or backend, if that is required by configuration constraints, but such a setup would not take full advantage of failover characteristics.

Figure 1–4 Facility Symbol




VM-0822A-AI

A facility name is mapped to specific physical nodes and their roles using the `CREATE FACILITY` command.

Figure 1–5 shows the logical relationship between client application, server application, frontends (FEs), routers (TRs), and backends (BEs) in the RTR environment at a specific location. The database is represented by the cylinder. Two facilities are shown (indicated by the large double-headed arrows), the User Accounts Facility and the General Ledger Facility. The User Accounts Facility uses three nodes, FE, TR, and BE, while the General Ledger Facility uses only two, TR and BE in the configuration shown. Its FEs are on nodes not shown in the figure, at another location.

Figure 1–5 Components in the RTR Environment



VM-0823A-AI

Clients send messages to servers to ask that a piece of work be done. Such requests may be bundled together into transactions. An RTR transaction consists of one or more messages that have been grouped together by a client application, so that the work done as a result of each message can be undone completely, if some part of that work cannot be done. If the system fails or is disconnected before all parts of the transaction are done, then the transaction remains incomplete.

Transaction

A **transaction** is a piece of work or group of operations that must be executed together to perform a consistent transformation of data. This group of operations can be distributed across many nodes serving multiple databases. Applications use services that RTR provides.

Because typically a transaction consists of several operations, a system or network failure at any step in the process will cause the transaction to be in doubt. RTR ensures that all transactions have the ACID properties so that all transactions are in a known state. (See the description of “Transactional messaging” for further clarification of transaction integrity.)

Transaction Controller

With the C++ API, the **Transaction Controller** manages transactions (one at a time), channels, messages, and events.

Transactional messaging

RTR provides **transactional messaging** in which transactions are enclosed in messages controlled by RTR.

Transactional messaging ensures that each transaction is complete, and not partially recorded. For example, a transaction or business exchange in a bank account might be to move money from a checking account to a savings account. The complete transaction is to remove the money from the checking account and add it to the savings account.

A transaction that transfers funds from one account to another consists of two individual updates: one to debit the first account, and one to credit the second account. The transaction is not complete until both actions are done. If a system performing this work goes down after the money has been debited from the checking account but before it has been credited to the savings account, the transaction is incomplete.

RTR Terminology

With transactional messaging, RTR ensures that a transaction is “all or nothing”—either fully completed or discarded; either both the checking account debit and the savings account credit are done, or the checking account debit is backed out and not recorded in the database. RTR transactions have the ACID properties.

Nontransactional messaging

An application will also contain nontransactional tasks such as writing diagnostic trace messages or sending a *broadcast* message about a change in a stock price after a transaction has been completed.

Transaction ID


Every transaction is identified on initiation with a transaction identifier or *transaction ID* (TID), with which it can be logged and tracked. RTR guarantees that TIDs are unique.

To reinforce the use of these terms in the RTR context, this section briefly reviews other uses of configuration terminology.


Tiers

A traditional two-**tier** client/server environment is based on hardware that separates application presentation and business logic (the clients) from database server activities. The client hardware runs presentation and business logic software, and server hardware runs database or data manager (DM) software, also called resource managers (RM). This type of configuration is illustrated in Figure 1–6. (In all diagrams, all lines are actually bidirectional, even when represented otherwise for clarity. For a given transaction, the initial action is typically from left to right.) In Figure 1–6, Application Presentation and Business Logic are the first tier, and the Database Server is the second tier.

Further separation into three tiers is achieved by separating presentation software from business logic on two systems, and retaining a third physical system for interaction with the database. This is illustrated in Figure 1–7, where Presentation and User Interface are the first tier, the Application Server and Business Logic are the second tier, and the Database Server is the third tier.

Figure 1–6 Two-Tier Client/Server Environment

VM-0824A-AI

Figure 1–7 Three-Tier Client/Server Environment

VM-0825A-AI

RTR extends the three-tier model, which is based on hardware, to a multilayer, or multicomponent software model.

RTR Software Components


RTR provides a multicomponent software model where client applications running on frontends, combined with routers and server applications running on backends cooperate to provide reliable service and transactional integrity. Application users interact with the client (presentation layer) on the node with the frontend role that forwards messages to the current router. The router in turn routes the messages to the current, appropriate backend, where server applications reside, for processing. The

connection to the current router is maintained until the current router fails or connections to it are lost.

All RTR software components can reside on a single node but are typically deployed on different nodes to achieve modularity, scalability, and redundancy for availability. During initial application development, it can be convenient to use a single physical node for all RTR roles and application software.

With different physical systems, if one node goes down or off line, another router or backend node can take over application processing. In a slightly different configuration, you could have an application that uses an external applet component running on a browser that connects to a client running on the RTR frontend. Such a configuration is shown in Figure 1–8. In the figure, the applet is separate from but connects to the client application written to work directly with RTR.

Figure 1–8 Browser Applet Configuration



VM-0826A-AI


The RTR client application could be an ASP (Active Server Page) or a process interfacing to the webserver through a standard interface such as CGI (Common Gateway Interface) script.

RTR provides automatic software failure tolerance and failure recovery in multinode environments by sustaining transaction integrity in spite of hardware, communications, application, or site failures. This automatic failover and recovery of service can exploit redundant or underutilized hardware and network links.

For example, you could use an underutilized system as a standby server in certain configurations.

As you modularize your application and distribute its components on frontends and backends, you can identify usage bottlenecks, add new nodes, and provide redundancy to increase availability. Adding backend nodes can help divide the transactional load and distribute it more evenly. For example, you could have a single node configuration as shown in Figure 1–9, RTR with Browser, Single Node, and Database. A single node configuration can be useful during development, but would not normally be used when your application is deployed.


Figure 1–9 RTR with Browser, Single Node, and Database



VM-0827A-AI

When applications are deployed, often the frontend is separated from the backend and router, as shown in Figure 1–10.

Figure 1–10 RTR Deployed on Two Nodes




VM-0828A-AI

RTR Terminology

In this example, the frontend with the client application resides on one node, and the router with the server application reside a node that has both the router and backend roles. This is a typical configuration where routers are placed on backends rather than on frontends. A further separation of workload onto three nodes is shown in Figure 1–11. However, in this configuration, there remain several single points of failure where one node/role or a network outage can disrupt processing of transactions.

Figure 1–11 RTR Deployed on Three Nodes



VM-0829A-AI

While this three-node configuration separates transaction load onto three nodes, it does not provide for continuing work if one of the nodes fails or becomes disconnected from the others. In many applications, there is a need to ensure that there is a server always available to access the database.


Standby server

In this case, a **standby server** will do the job. A standby server (see Figure 1–12) is a process or application that can take over when the primary server is not available, due to hardware failure, application software failure or network outage.

Both the primary and the standby server have the capability to access the same database, but the primary processes all transactions unless it is unavailable. On the other hand, the standby processes transactions only when the primary becomes unavailable. When not being used to process transactions, the standby CPU can do other work.

The standby server is usually placed on a node other than the node where the primary server runs, and should be, to avoid being a single point of failure. Network capability, clustering or disk-sharing technology, and appropriate software must be available on both primary and standby backend systems when running RTR.

Figure 1–12 Standby Server Configuration




VM-0830A-AI

Shadow Server and Transactional shadowing

To increase transaction processing availability, transactions can be shadowed with a **shadow server**, as shown in Figure 1–13. The system where the shadow server runs can be made available with clustering technology. A shadow server eliminates the single point of failure that is evident in Figure 1–12. In a shadow configuration, the second database of Figure 1–13 is available even when the first is not.

Use of a shadow server is called **transactional shadowing** and is accomplished by having a second location, often at a different site, where transactions are also recorded. Data are recorded in two separate data stores or databases. The router knows about both backends and sends all transactions to both backends. RTR provides the server application with the necessary information to keep the two databases synchronized.

Figure 1–13 Transactional Shadowing Configuration

VM-0831A-AI

In the RTR environment, one data store (database or data file) is elected the primary, and a second data store is made the shadow. The shadow data store is a copy of the data store kept on the primary. If either data store becomes unavailable, all transactions continue to be processed and stored on the surviving data store. At the same time, RTR makes a record of (remembers) all transactions stored only on the shadow data store in the **RTR journal** by the shadow server.

When creating the configuration used by an application and defining the nodes where a facility has its frontends, routers, and backends, the setup must also define which nodes will have journal files. Each backend in an RTR configuration must have a journal file to capture transactions when other nodes are unavailable. When the primary server and data store become available again, RTR replays the transactions in the journal to the primary data store through the primary server. This brings the data store back into synchronization.

With transactional shadowing, there is no requirement that hardware, the data store, or the operating system at different sites be the same. You could, for example, have one site running OpenVMS and another running Windows; the RTR transactional commit process would be the same at each site. Because the database resides at both sites, either backend can have an outage and all transactions will still be processed and recovered.

Note

Transactional shadowing shadows only transactions controlled by RTR.

For full redundancy to assure maximum availability, a configuration could employ **disk shadowing** in clusters at separate sites coupled with **transactional shadowing** across sites. Disk shadowing used in specific cluster environments copies data to another disk to ensure data availability. Transactional shadowing copies only transactional data.

Additionally, an RTR configuration typically deploys multiple frontends running client applications with connections to several routers to ensure continuing operation if a particular router fails.

RTR Server Types

In the RTR environment, in addition to the placement of frontends, routers, and servers, the application designer must determine what server capabilities to use. RTR provides four types of software servers for application use:

- Standby servers
- Transactional shadow servers
- Concurrent servers
- Callout servers

These are described in the next few paragraphs. You specify server types to your application in RTR API calls.

RTR server types help to provide continuous availability and a secure transactional environment.

RTR Server Types

Standby server

The **standby server** remains idle while the RTR primary backend server performs its work, accepting transactions and updating the database. When the primary server fails, the standby server takes over, recovers any in-progress transactions, updates the database, and communicates with clients until the primary server returns. There can be many instances of a standby server. Activation of the standby server is transparent to the user.

A typical standby configuration is shown in Figure 1–12, Standby Server Configuration. Both physical servers running the RTR backend software are assumed by RTR to connect to the same database. The primary server is typically in use, and the standby server can be either idle or used for other applications, or data partitions, or facilities. When the primary server becomes unavailable, the standby server takes over and completes transactions as shown by the dashed line. Primary server failure could be caused by server process failure or backend (node) failure.

Standby in a cluster

The intended and most common use of a standby server is in a recognized cluster environment. In a noncluster or unrecognized cluster environment, seamless failover of standbys is not guaranteed. For RTR, clusters supported by OpenVMS and Tru64 UNIX are recognized clusters, whose processing is controlled by a lock manager. Windows and Sun clusters can use disk-sharing, unrecognized cluster technology.

Standby servers are “spare” servers that automatically take over from the main backend if it fails. This takeover is transparent to the application.

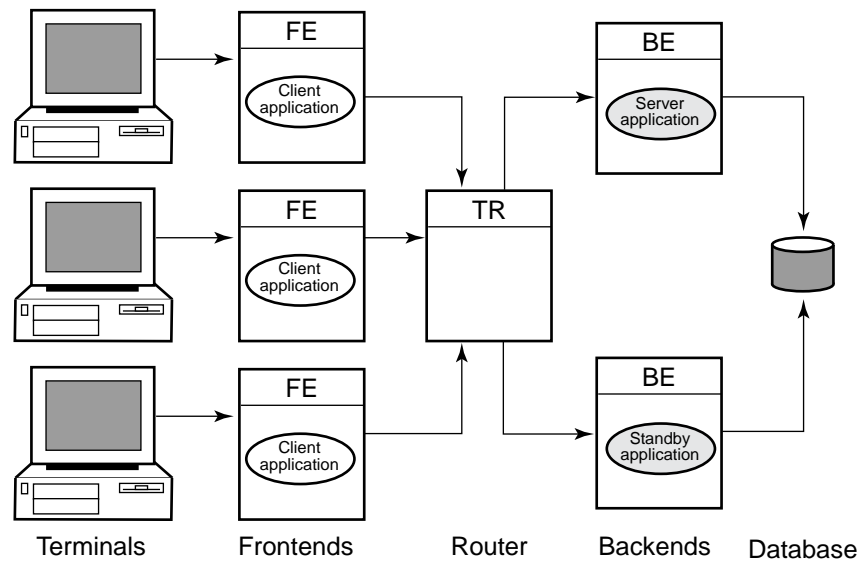
Figure 1–14 shows a simple standby configuration. The two backend nodes are members of a cluster environment, and are both able to access the database.

For any one key range, the main or primary server (Server application) runs on one node while the standby server (Standby application) runs on the other node. The standby server process is running, but RTR does not pass any transactions to it. Should the primary node fail, RTR starts passing transactions to the Standby application.

RTR Server Types

Note that one node can contain the primary servers for one key range and standby servers for another key range to balance the load across systems. This allows the nodes in a cluster environment to act as standby for other nodes without having idle hardware. When setting up a standby server, both servers must have access to the same journal.

Figure 1–14 Standby Servers



VM-0833A-AI

Transactional shadow server

The **transactional shadow server** places all transactions recorded on the primary server on a second database. The transactional shadow server can be at the same site or at a different site, with networking capability available.

When one member of a shadow set fails, RTR remembers the transactions executed at the surviving site in a journal, and replays them when the failed site returns. Only after all journaled transactions are recovered does the recovering site fully process new online transactions. During recovery, new transactions are processed at the surviving site and added to the journal for the recovering site.

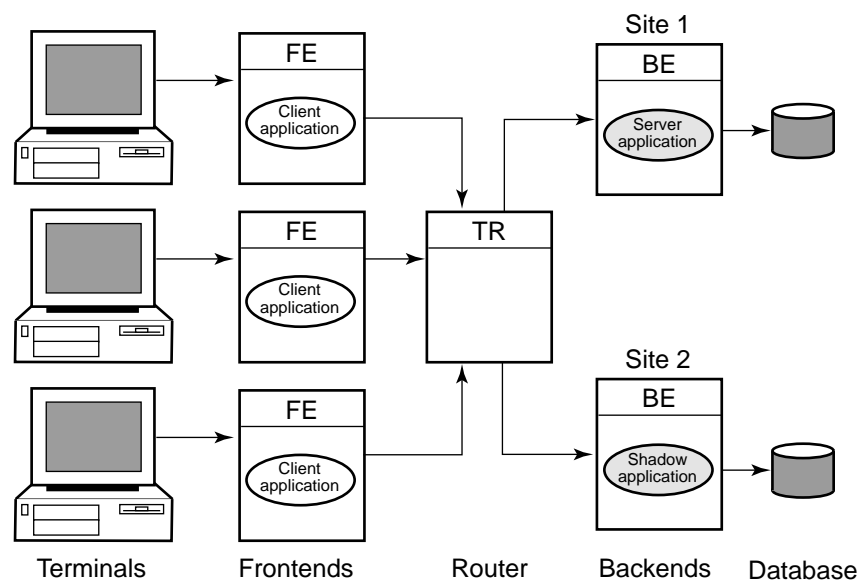
RTR Server Types

Transactional shadowing is done by partition. A transactional shadow configuration can have only two members of the shadow set.

Shadow servers are servers on separate backends that handle the same transactions in parallel on identical copies of the database.

Figure 1–15 shows a simple shadow configuration. The main backend server application at Site 1 and the shadow server (Shadow application) at Site 2 both receive every transaction for the data partition they are servicing. Should Site 1 fail, Site 2 continues to operate without interruption. Sites can be geographically remote, for example, available at separate locations in a wide area network (WAN).

Figure 1–15 Shadow Servers




VM-0834A-AI

Concurrent server

The **concurrent server** is an additional instance of a server application running on the same node. RTR delivers transactions to a free server from the pool of concurrent servers. If one server fails, the transaction in process is replayed to another server in the concurrent pool. Concurrent servers are designed primarily to increase throughput and can exploit Symmetric Multiprocessing (SMP) systems. Figure 1–16, Concurrent Servers, illustrates the use of concurrent servers sending transactions to the same partition on a backend, the partition A-N.

Concurrent servers allow transactions to be processed in parallel to increase throughput. Concurrent servers deal with the same database partition, and may be implemented as multiple channels within a single process or as one channel in separate processes. The application designer must determine if transactions can be processed concurrently by the database or server application. Deadlocks can occur if every transaction competes for the same database lock outside the RTR server application.

Figure 1–16 Concurrent Servers



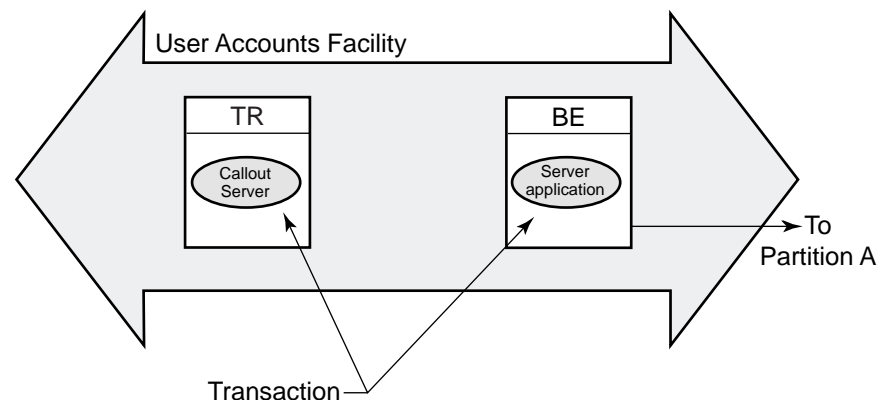
RTR Server Types

Callout server

The **callout server** enables message authentication on transaction requests made in a given facility, and could be used, for example, to provide audit trail logging. A callout server can run on either backend or router nodes. A callout server receives a copy of all messages in a facility. Because the callout server votes on the outcome of each transaction it receives, it can veto any transaction that does not pass its checks.

A callout server is facility-based, not partition-based; any message arriving at the facility is routed to both the server and the callout. A callout server is enabled when the facility is defined. Figure 1–17 illustrates the use of a callout server that authenticates every transaction in a facility.

Figure 1–17 A Callout Server



VM-0836A-AI

To authenticate any part of a transaction, the callout server must vote on the transaction, but does not write to the database. RTR does not replay a transaction that is only authenticated.

RTR Server Types

Authentication

RTR callout servers provide facility-based processing for authentication. For example, a callout server can enable verification checks to be carried out on all requests in a given facility.

Callout servers run on backend or router nodes. They receive a copy of every transaction either delivered to or passing through the facility.

Callout servers offer the following advantages:

- The checking code is completely separated from other application code. The application does not have to be modified to add an authentication check or logging.
- The check can run in parallel with the database updates; this can improve response times.
- The check can be run on the router hardware.

Because this technique relies on backing out unauthorized transactions, it is most suitable when only a small proportion of transactions are expected to fail the check, so as not to have a performance impact.

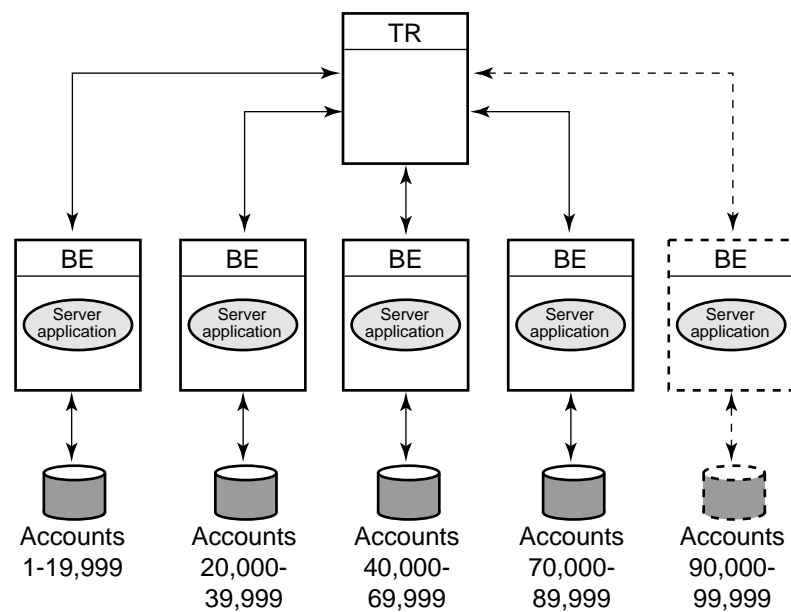
Partition

When working with database systems, partitioning the database can be essential to ensuring smooth and untrammelled performance with a minimum of bottlenecks. When you **partition** your database, you locate different parts of your database on different disk drives to spread both the physical storage of your database onto different physical media and to balance access traffic across different disk controllers and drives.

For example, in a banking environment, you could partition your database by account number, as shown in Figure 1–18. A partition is a part or segment of your database.

Key range

Once you have decided to partition your database, you use key ranges in your application to specify how to route transactions to the appropriate database partition. A **key range** is the range of data held in each partition. For example, the key range for the first partition in the bank partitioning example goes from 00001 to 19999.

Figure 1–18 Bank Partitioning Example

VM-0837A-AI

You can assign a partition name in your application program or have it set by the system manager. Note that sometimes the terms key range and partition are used as synonyms in code examples and samples with RTR, but strictly speaking, the key range defines the partition. A partition has both a name, its partition name, and an identifier generated by RTR — the partition ID. The properties of a partition (callout, standby, shadow, concurrent, key segment range) can be defined by the system manager with a `CREATE PARTITION` command. For details of the command syntax, refer to the *HP Reliable Transaction Router System Manager's Manual*.

A significant advantage of the partitioning shown in the bank example is that you can add more account numbers without making changes to your application; you need only add another server and disk drive for the new account numbers. For example, say you need to add account numbers from 90,000 to 99,999 to the basic configuration of Figure 1–18, Bank Partitioning Example. You can add these accounts and bring them on line

easily. The system manager can change the key range with a command, for example, in an overnight operation, or you can plan to do this during scheduled maintenance.


A partition can also have multiple standby servers.

Standby Server Configurations

A node can be configured as a primary server for one key range and as a standby server for another key range. This helps to distribute the work of the standby servers. Figure 1–19 illustrates this use of standbys with distributed partitioning.

As shown in Figure 1–19, the Server application on backend A is the primary server for accounts 1 to 19,999 and the Server application on backend B is the standby for these same accounts, shown as a dashed line. The Server application on backend B is the primary for accounts 20,000 to 39,999 and the Server application on backend A can be the standby for these same accounts (not shown in the figure). For clarity, account numbers are shown only for the primary servers and one standby server.

Figure 1–19 Standby with Partitioning



VM-0838A-AI

Anonymous clients

RTR supports anonymous clients; that is, clients can be set up in a configuration using wildcarded node names.

Tunnel

RTR can also be used with firewall tunneling software, which supports secure internet communication for an RTR connection, either client-to-router, or router-to-backend.

RTR Networking Capabilities

Depending on operating system, RTR uses TCP/IP or DECnet as underlying transports for the virtual network (RTR facilities) and can be deployed in both local area and wide area networks. PATHWORKS 32 is required for DECnet configurations on Windows.

Architectural Concepts

This chapter introduces concepts on basic transaction processing and RTR architecture.

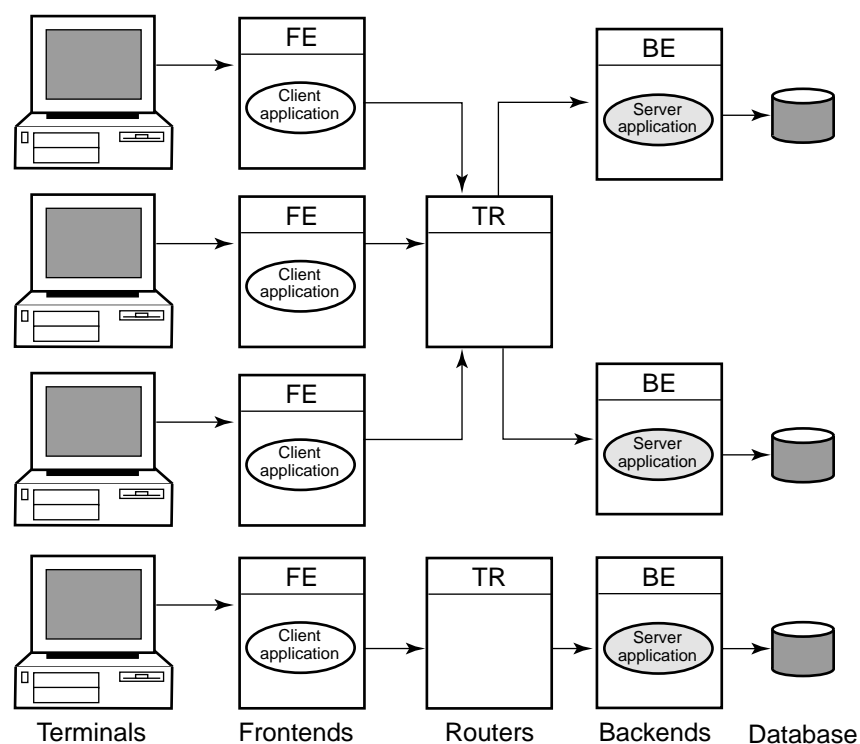
The Three-Tier Architecture

RTR is based on a three-tier physical architecture consisting of frontend (FE) roles, backend (BE) roles and router (TR) roles. The roles are shown in Figure 2–1. (In this and subsequent diagrams, rectangles represent physical nodes, ovals represent application software, and cylinders represent the disks storing the database. The nodes connected to the actual database usually run the database software that controls the database.)

In addition to the physical configuration where RTR is deployed, software plays a critical part, extending the tier concept to more than three tiers. On certain pieces of hardware, client application software runs, and on others, server application software runs. Users can connect to nodes that are running the frontend role with appropriate non-RTR software. For example, a user can have a PC where RTR runs; in this case, the PC has the frontend role. Or a user could use a PC running, say Pathworks, to connect to another node that has the frontend role and run the RTR client application from there. This would be a multitier configuration.

The Three-Tier Architecture

Figure 2-1 The Multitier Model



VM-0839A-AI

Client application processes run on nodes defined to have the frontend role. This tier allows computing power to be provided locally at the end-user site for transaction acquisition and presentation.

Server application processes (represented by “Server application” in Figure 2-1) run on nodes defined to have the backend role. This tier:

- Enables the database to be distributed geographically
- Permits replication of servers to cope with network, node or site failures
- Allows computer resources to be added to meet performance requirements

The Three-Tier Architecture

- Allows performance or geographic expansion while protecting the investments made in existing hardware and application software

The router tier contains no application software unless running callout servers. This tier reduces the number of logical network links required on frontend and backend nodes and helps ensure good performance even in an unstable network. It also decouples the backend tier from the frontend tier so that configuration changes in the (frequently changing) user environment have little influence on the transaction processing and database (backend) environment.

The three-tier model can be mapped to any system topology. More than one role may be assigned to any particular node. For example, on a system with few frontends, the router and backend tiers can be combined in the same nodes. During application development and test, all three roles can be combined in one node.

The nodes used by an application and their configuration roles are specified using RTR configuration commands. RTR lets application code be completely location and configuration independent.

RTR Facilities Bridge the Gap

Many applications can use RTR at the same time without interfering with one another. This is achieved by defining a separate facility for each application. A facility can be thought of as an application network.

For example, when an application calls the C++ API `RTRFacilityManager` to manage a channel or the C API `rtr_open_channel()` routine to declare a channel as a client or server, it specifies the name of the facility it will use.

Refer to the *HP Reliable Transaction Router System Manager's Manual* for information on how to define facilities.

Broadcasts

Sometimes an application has a requirement to send unsolicited messages to multiple recipients.

An example of such an application is a commodity trading system, where the clients submit orders and also need to be informed of the latest price changes.

The RTR broadcast capability meets this requirement.

Recipients subscribe to a class of broadcasts; a sender broadcasts a message in this class, and all interested recipients receive the message. However, broadcast reception is not guaranteed; network or node outages can cause a particular client to fail to receive a broadcast message.

RTR permits clients to broadcast messages to one or more servers, or servers to broadcast to one or more clients. If a server needs to broadcast a message to another server, it must open a second channel as a client.

Flexibility and Growth

RTR allows you to cope easily with changes in:

- Network demand
- User access patterns
- Volume of data

Because an RTR-based system can be built using multiple systems at each functional tier, it easily lends itself to step-by-step growth, avoiding unused capacity at each stage. With your system still up and running, it is possible to:

- Create and delete concurrent server processes.
- Add or remove nodes (frontend, router or backend).

This means you do not need to provide spare capacity to allow for growth.

RTR also allows parallel execution. This means that different parts of a single transaction can be processed in parallel by multiple servers.

RTR provides a comprehensive set of monitoring tools to help you evaluate the volume of traffic passing through the system. These tools can help you respond to unexpected load changes, making you aware of system degradation that you can sometimes alleviate by altering the system configuration dynamically.

Transaction Integrity

RTR greatly simplifies the design and coding of distributed applications because, with RTR, database actions can be bundled together into transactions.

To ensure that your application deals with transactions correctly, its transactions must have the ACID properties, fundamental properties of transaction processing systems. A transaction that has the ACID properties is:

- Atomic
- Consistent
- Isolated
- Durable

An *atomic* transaction is all or nothing; that is, either the entire transaction is totally committed or totally rolled back. A *consistent* transaction either creates a new, valid state of data, or, from any failure, returns all data to its state as it was before the start of the transaction. An *isolated* transaction does not cause changes to shared resources until commitment of the transaction. A *durable* transaction survives system and media failures after transaction commitment. A durable transaction is thus both persistent and stable.

For more detail on these properties and their use in transaction processing, refer to the *HP Reliable Transaction Router Application Design Guide*.

The Partitioned Data Model

One goal in designing for high transaction throughput is reducing the time that users must wait for shared resources.

While many elements of a transaction processing system can be duplicated, one resource that must be shared is the database. Users compete for a shared database in three ways:

- For use of the disk
- For locks on database records
- For the CPU resources needed to access the database

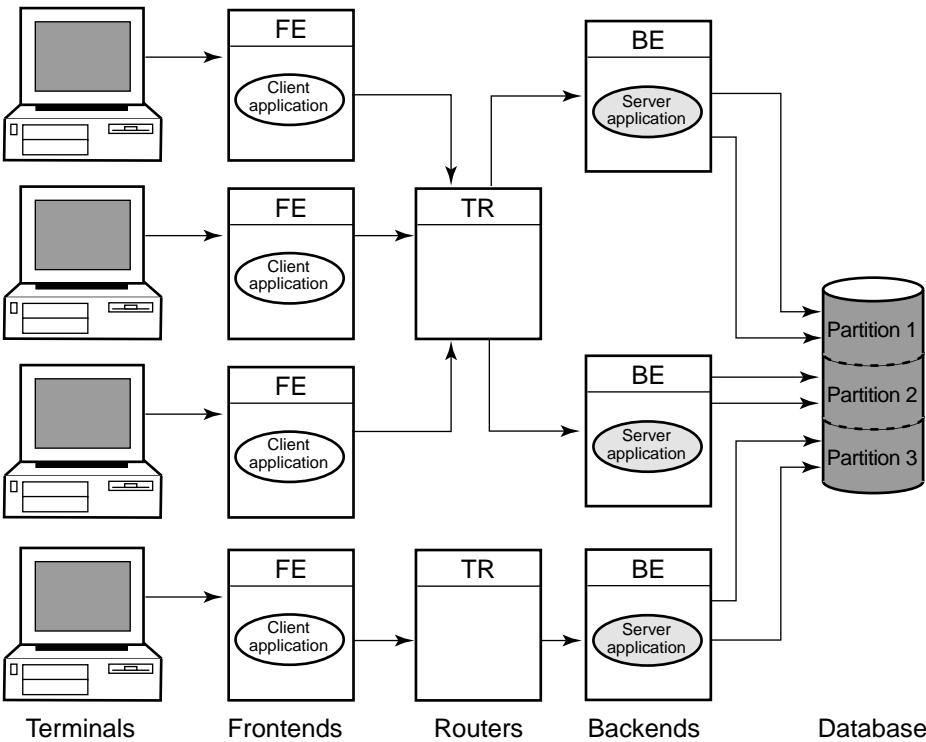
This competition can be alleviated by spreading the database across several backend nodes, where each node is responsible for a subset of the data contained in a partition. RTR enables you to implement this partitioned data model, shown roughly in Figure 2–2 where the database has three partitions. RTR routes messages to the correct partition on the basis of an application-defined key. For a more complete description of partitioning as provided with RTR, refer to the *HP Reliable Transaction Router Application Design Guide*.

Each RTR API provides the capability to use partitions. For specific information on declaring and using partitions, see the RTR documentation for the system manager and the applications programmer.

Object-Oriented Programming

Java objects and the RTR C++ foundation classes map traditional RTR functional programming routines into object-oriented programming models. Using the power and features of the Java objects or the foundation classes requires understanding of the differences between functional and object-oriented programming concepts. Table 2–1 compares the worlds of functional programming and object-oriented programming.

Figure 2-2 Partitioned Data Model



VM-0840A-AI

Table 2–1 Functional and Object-Oriented Programming Compared

Functional Programming	Object-Oriented Programming
A program consists of data structures and algorithms.	A program consists of a team of cooperating objects.
The basic programming unit is the function, that when run, implements an algorithm.	The basic programming unit is the class, that when instantiated, implements an object.
Functions operate on elemental data types or data structures.	Objects communicate by sending messages.
An application's architecture consists of a hierarchy of functions and sub-functions.	An applications architecture consists of objects that model entities of the problem domain. Objects' relationships can vary.

Objects

In the object-oriented environment, a program or application is a grouping of cooperating objects. The basic programming unit is the class. Instantiating, or declaring, an instance of, a class implements an object. RTR provides object-oriented programming capabilities with the C++ API, described in the *HP Reliable Transaction Router C++ Foundation Classes* manual and the Java API, described in the *JRTR Getting Started* manual. Objects are instances of a class. In a transaction class, each transaction is an object. An object is an instantiated (declared) class. Its state and behavior are determined by the attributes and methods defined in the class. An object or class is defined by its:

- State (attributes)
- Behavior (methods)
- Identity (name at instantiation)

The name given at object declaration is its identity. In Example 2–1, the two dog objects King and Fifi are instances of Dog. The Dog class is declared in a header (Dog.h) file and implemented in a .cpp file.

Example 2–1 Objects-Defined Sample

```

Dog.h:
class Dog
{ ...
};
main.cpp:
#include "Dog.h"
main()
{
    Dog King;
    Dog Fifi;
}

```

Messages

Objects communicate by sending messages. This is done by calling an object's methods.

Some principal categories of messages are:

- Constructors: Create objects
- Destructors: Delete objects
- Selectors: Return part or all of an object's state. For example, a Get method
- Modifiers: Change part or all of an object's state. For example, a Set method
- Iterators: Access multiple element objects within a container object. For example, an array.

Class Relationships

Classes can be related in the following ways:

- Simple association: One class is aware of another class. For example, "Dog object is associated with a Master object." This is a "Knows a" relationship.
- Composition: One class contains another class as part of its attributes. For example, "Dog objects contains Leg objects." This is a "Has a" relationship.
- Inheritance: A child class is derived from one or more parent, or base, classes. For example, "Mutt object derives from Collie object and Boxer object which both derive from Dog object." This is an "Is a" relationship. Inheritance enables the use of polymorphism.

Polymorphism

Polymorphism is the ability of objects, inherited from a common base or parent class, to respond differently to the same message. This is done by defining different implementations of the same method name within the individual child class definitions. For example: A DogArray object, "DogArray OurDogs[2];" refers to two element objects of class Dog, the base class:

- King, of class Doberman, is a derived or child class of Dog.
- Fifi, of class Minipoodle, is a derived or child class of Dog.

If, in a program, OurDogs[n]->Bark() is called in a loop, then:

- In iteration one ([1]), method King::Bark() is called.
- In iteration two ([2]), method Fifi::Bark() is called.

King's bark does not sound like Fifi's bark because each Bark() call is a separately defined method within its child object definition. The virtual parent class (Dog) method Bark() is defined in the class definition of Dog.

Object Implementation Benefits

The benefits of creating RTR solutions with objects include the following:

- Each major RTR concept is represented by its own individual foundation class.
- Simple methods within RTR classes transform features of RTR for streamlined solutions.
- Major classes include Get and Set methods for changing transaction states.
- Default handling code is provided for all Messages and Events, where appropriate.
- You do not need to provide handling code for all messages and events.
- The sending and receiving of data is abstracted to a higher level with transaction controller and data classes.
- No buffers and links coding is needed.
- Internal RTR information is accessible without a need to know RTR internals.

Java Support

RTR clients and servers can be Java applications that obtain the benefits of high availability, fault tolerance and scalability provided by RTR. RTR clients and servers employing Java technology use standard Java and J2EE interfaces for transaction management, data input/output, and database access.

For additional information, see the *JRTR Getting Started* manual and associated online documentation.

XA Support

Within its C API, RTR provides the capability of using the XA interface to work with XA-compliant database systems. The XA interface is part of the X/Open DTP (Distributed Transaction Processing) standard. It defines the interface that transaction managers (TM) and resource managers (RM) use to perform the two-phase commit protocol. (Resource managers are underlying database systems such as ORACLE RDBMS, Microsoft SQL Server, and others.) This interface is used by TM-to-RM exchanges to coordinate a transaction from within an application program.

If your database application supports XA, you have less to implement in your application environment; use of XA can also increase the portability of your application.

For details on using XA with RTR, refer to the *HP Reliable Transaction Router C Application Programmer's Reference Manual* and the *HP Reliable Transaction Router Application Design Guide*.

3

Reliability Features

This chapter addresses:

- RTR server types
- Failover and recovery
- Recovery scenarios

RTR Server Types

Reliability in RTR is enhanced by the use of:

- Concurrent servers
- Standby servers
- Shadow servers
- Callout servers
- Router failover

Note that, conceptually, servers can be contrasted as follows:

- Concurrent servers handle *similar* transactions which access the same data partition and run on the same node.

When transaction throughput is constrained by your server application, consider adding a second instance of your server application with a concurrent server.

- Shadow servers handle the *same* transactions and run on different nodes.

RTR Server Types

When there is concern that your database is a single point of failure, add a shadow server, if possible at a different physical location.

- Standby servers provide a node that can take over processing on a data partition when the primary server or node fails.

When there is concern that your server application or the node where it is running is a single point of failure in your configuration, configure a standby server to be ready to take over.

- Callout servers run on backends or routers and receive all messages within a facility so that authentication and logging operations can be performed in parallel.

Use a callout server to add processing logic (authentication or logging) to your transactions without modifying your server application.

All servers are further described in the earlier section on RTR Terminology.

Failover and Recovery

RTR provides several capabilities to ensure failover and recovery under several circumstances.

Router Failover

Frontend nodes automatically connect to another router if the one being used fails. This reconnection is transparent to the application.

Routers are responsible for coordinating the two-phase commit for transactions. If the original router coordinating a transaction fails, backend nodes select another router that can ensure correct transaction completion.

**Backend
Restart
Recovery**

Transactions in the process of being committed at the time of a failure are recovered from RTR's disk journal. Recovery could be with a concurrent server, a standby server, or a restarted server created when the failed backend restarts.

Correct ordering of the execution of transactions against the database is maintained.

**Transaction
Message
Replay**

Transaction messages that are lost in transit are resent when possible. The frontend and backend nodes keep an in-memory copy of all active messages for this purpose.

**Link Failure
Recovery**

In the event of a communications failure, RTR tries to reconnect the link or links until it succeeds.

Recovery Scenarios

This section describes how RTR recovers from different hardware and software failure. For additional information on failure and recovery scenarios, refer to the *HP Reliable Transaction Router Application Design Guide*.

**Backend
Recovery**

If standby or shadow servers are available on another backend node, operation of the rest of the system will continue without interruption, using the standby or shadow server.

If a backend processor is lost, any transactions in progress are remembered by RTR and later recovered, either when the backend restarts, or by a standby if one is present. Thus, the distributed database is brought back to a transaction-consistent state.

**Router
Recovery**

If a router fails and another router node is available, all in-progress transactions are transparently rerouted by the other router. System operation will continue without interruption.

Recovery Scenarios

Frontend Recovery

If a frontend is lost:

- All transactions committed but not completed on the frontend node at the time of failure will be completed.
- All transactions started but not committed on the frontend node at the time of failure will be aborted.

RTR Interfaces

Management Interfaces

RTR provides interfaces for system management (the management interfaces) and for development of transaction processing and management applications (the programming or application development interfaces).

The management interfaces are:

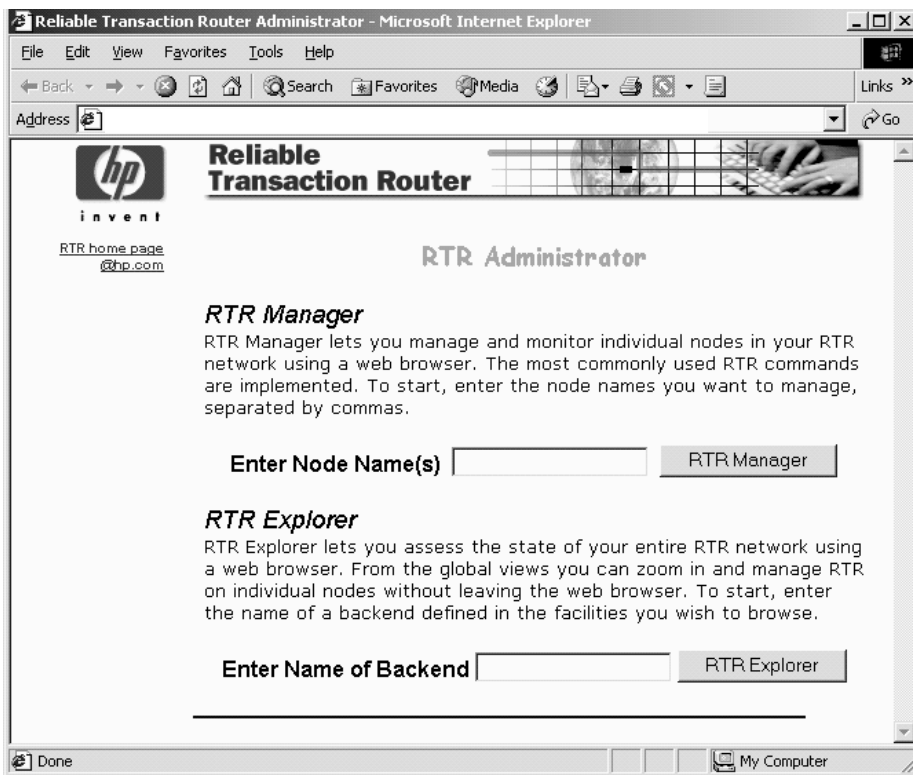
- The RTR Administrator, a browser interface. This interface includes:
 - RTR Manager
 - RTR Explorer
- The command line interface or CLI

The RTR Administrator lets you manage RTR, its facilities, nodes, and network links, with a point-and-click interface. It contains extensive help, both as inline popups, as linked help, and as links to current information. For example, inline popups describe short headings more fully, and the system manager can view many types of status as RTR and the applications under its control run.

Figure 4–1 shows the RTR Administrator screen where you select whether to use the RTR Manager or the RTR Explorer.

RTR Interfaces

Figure 4–1 RTR Administrator

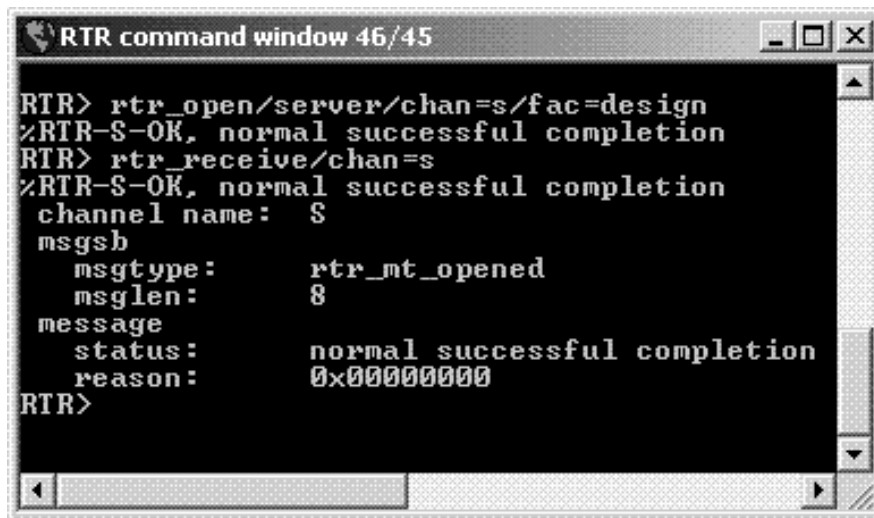


vm-1152A-AI

The RTR CLI contains all RTR system manager commands and calls to all RTR C API routines such as `rtr_open_channel` or `rtr_create_facility`. You can use either the RTR Manager or the RTR CLI to manage your RTR configuration. You can also use the command line interface to write short RTR C applications for testing and experimentation. The CLI is described in the *HP Reliable Transaction Router System Manager's Manual*. Its use is illustrated in this chapter.

Figure 4–2 shows the RTR command line interface.

Figure 4-2 RTR Command Line Interface



```

RTR command window 46/45
RTR> rtr_open/server/chan=s/fac=design
%RTR-S-OK, normal successful completion
RTR> rtr_receive/chan=s
%RTR-S-OK, normal successful completion
channel name:  S
msgsh
  msgtype:      rtr_mt_opened
  msglen:       8
message
  status:       normal successful completion
  reason:       0x00000000
RTR>

```

VM-1103a-AI

Programming Interfaces

RTR provides several programming or application development interfaces for design and implementation of transaction processing programs. They include the following:

- The object-oriented RTR Java interface
You can use this API for new development, and, where appropriate, for new development with existing applications. This API can be used to implement applications with RTR using Java and J2EE technologies.
- The object-oriented API for C++ programming
You can use this API for new development and, where appropriate, for new work with existing applications. An application can contain both object-oriented classes and Portable API calls. The C++ API can be used to implement both management and transaction processing applications on all platforms supported by RTR.
- The RTR API for C programming
This interface was the first multiplatform API available with RTR.

RTR Interfaces

- An interface that enables use of an X/Open Distributed Transaction Processing-conformant resource manager
This interface, invoked through the RTR management interfaces, enables RTR applications to be used with X/Open-compliant resource managers such as Oracle8.
- The OpenVMS API containing OpenVMS calls
This API, supported on OpenVMS only, is obsolete for new development. To take advantage of new RTR features and capabilities, such applications can be rewritten with one of the newer APIs. Older applications will continue to run with later versions of RTR.

The RTR application programming interfaces, where available, are identical on all hardware and operating system platforms that support RTR. The object-oriented C++ API is fully described in the *HP Reliable Transaction Router C++ Foundation Classes* manual. The C-programming API is fully described in the *HP Reliable Transaction Router C Application Programmer's Reference Manual*. Both APIs are used in examples in the *HP Reliable Transaction Router Application Design Guide*. The Java J2EE interface is described in the JTA material in the RTR JRTR kit. The XA interface is described in materials from X/Open.

Application Development

The transaction processing environment poses special challenges for the development of applications, challenges best addressed by following a defined methodology such as the software development life cycle or object-oriented design fleshed out with use cases.

The software development life cycle consists of the following phases that are to be viewed as iterative:

- Gathering requirements (what is needed?)
- Developing a high-level design (what will the transaction processing application do?)
- Constructing a detailed design (explicitly, what will each part of the application do and what are its intended results?)
- Coding and unit testing
- Integration and system testing/deployment
- Maintenance

Many books are available to assist the developer with both design and development, including:

- J. Gray, A. Reuter, *Transaction Processing: Concepts and Techniques*, Morgan Kaufmann, San Mateo, CA, 1992
- Philip A. Bernstein, Eric Newcomer, *Principles of Transaction Processing*, Morgan Kaufmann, San Francisco, CA, 1997

Object-oriented methods and practice are described and elaborated on in many books, including:

- James Rumbaugh, Michael Blaha, William Lorensen, *Object-Oriented Modeling and Design*, Prentice Hall, Englewood Cliffs, NJ, 1991
- Martin Fowler with Kendall Scott, *UML Distilled*, Addison-Wesley, Reading, MA 1997

Table 4–1 summarizes the RTR application development interfaces and their typical use.

Table 4–1 RTR Development Interfaces and their Use

With this interface:	You can write:
RTR Java interface	application programs
RTR C++ Foundation Classes	application programs system management programs
RTR C programming interface	application programs

RTR Management

You can manage RTR from several locations:

- from a node on which RTR is running
- from a remote node from which you send RTR commands to a node running RTR
- from a web browser that can be on or access a node running RTR

RTR Administrator

With the RTR Administrator, you have a network-browser-like display from which you can view RTR status and issue certain RTR commands with a point-and-click operation. The RTR Administrator contains both the RTR Manager, with which you set up and manage your RTR configuration, and the RTR Explorer, with which you observe and monitor your RTR configuration. Online help provided with the RTR Administrator includes popups for screen headings, popups for some RTR Explorer information, and extensive help for RTR commands.

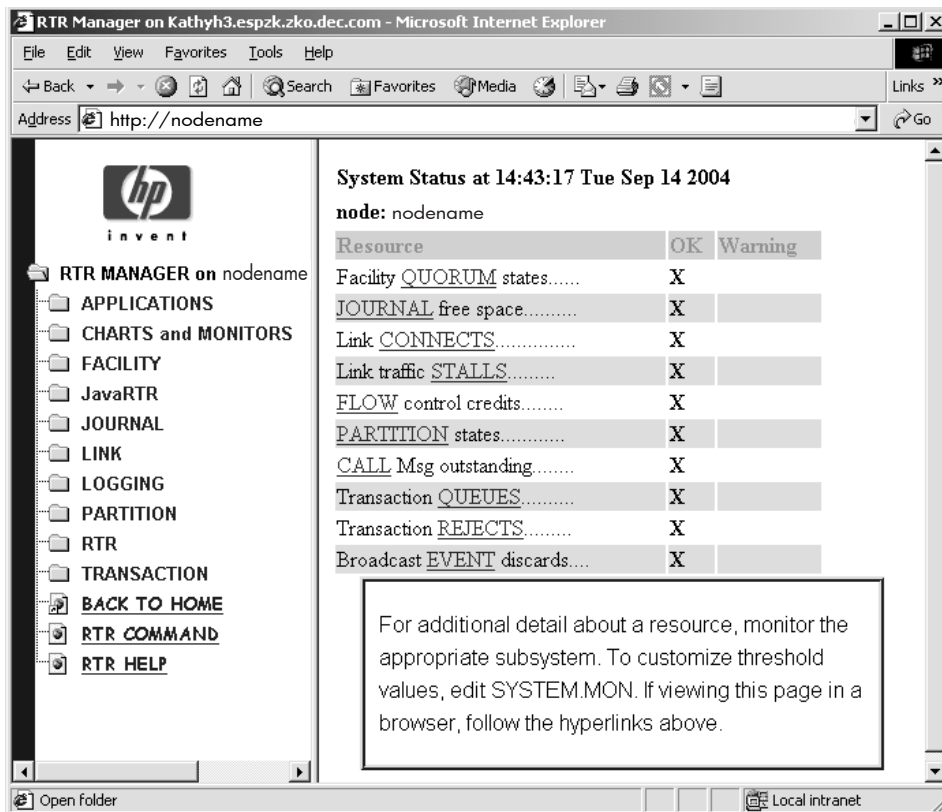
In the RTR Manager or RTR Explorer, buttons have the following effect:

This button:	Takes you to:
Back	The previous level
Back to Home	To the RTR Administrator

RTR Manager

Figure 4–3 shows the first RTR Manager screen through which you manage RTR and RTR applications. Not all RTR CLI commands are accessible from the RTR Manager RTR Command link; those rarely used are available only through the RTR CLI command window. The RTR Manager provides help for input screens, logging windows, and links between displays.

Figure 4–3 RTR Manager



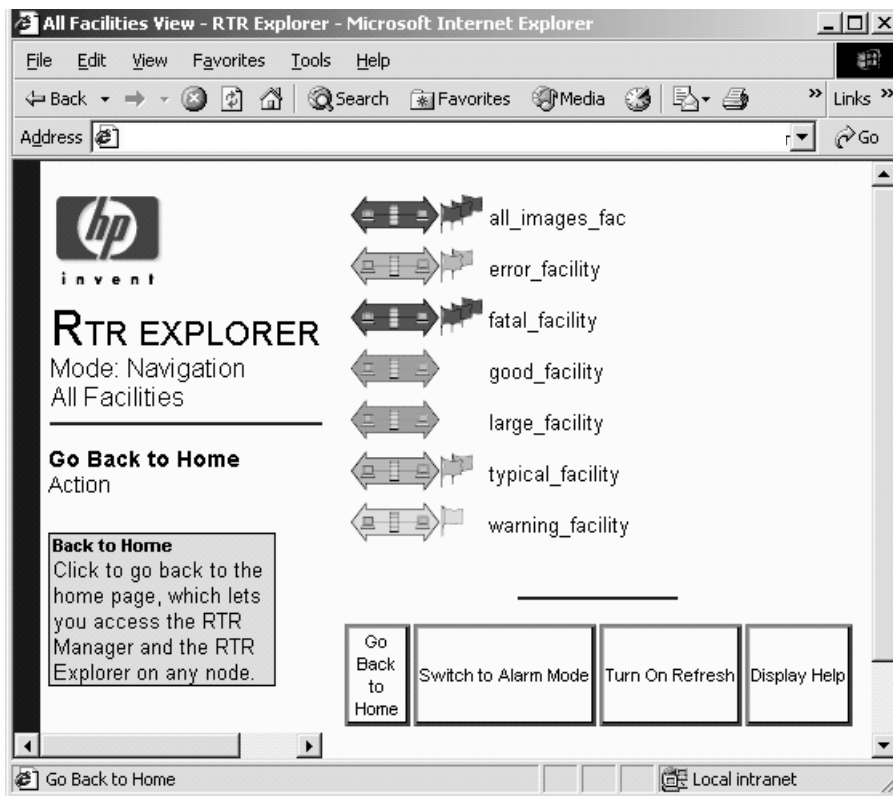
VM-1158A-AI

RTR Explorer

Figure 4–4 shows a sample screen of the RTR Explorer with several defined facilities. The RTR Explorer lets the system manager or administrator view the entire RTR configuration by facility and by node, and assess the state of the RTR network and the state of any individual node or facility in the network.

The All Facilities View (Figure 4–4) shows all facilities by name, with icons showing status of normal or one of three levels of alert (warning, error, or fatal). Additional views either by facility or node enable the manager to drill down, with a simple point-and-click, to the facility or node of interest.

Figure 4-4 RTR Explorer: View of All Facilities

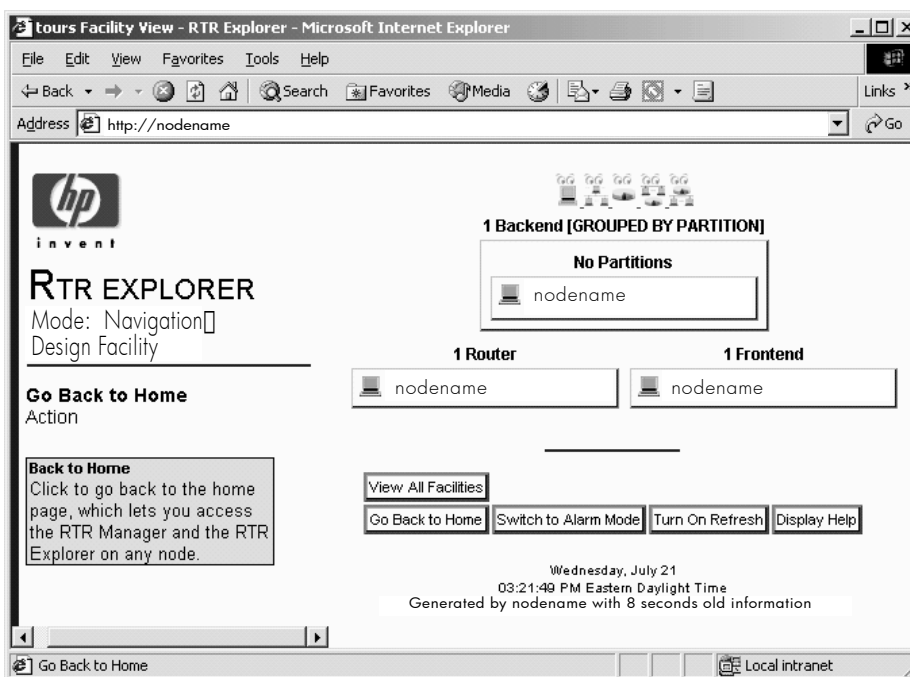


vm-1155A-AI

Information available for each facility includes the facility name, the alerts associated with nodes participating in the facility, and the state of the facility. Information for each node includes its name, role, cluster name, partition names, partition states, node state, and alerts.

Additional views either by facility or by node enable the administrator to zoom, with a simple point-and-click, to the facility or node of interest. From the node view, the administrator can open the RTR Manager for the node. Hence, RTR Explorer enables the administrator to both view and manage the state of every facility and node.

Figure 4–5 RTR Explorer: View of One Facility



VM-1151A-AI

Nodes can monitor themselves for alerts. Each alert can be set at progressive levels of severity – first *Warning*, second *Error*, and third *Fatal*. The severity of an alert indicates the urgency of the alert. *Warning* means RTR may or may not be operating normally, but something needs to be looked at. *Error* means that RTR is likely not operating normally, but may be able to continue operating. *Fatal* means that RTR cannot continue to operate unless the alert is resolved. See the REMEMBER EXPRESSION command in the System Manager’s Manual for how to define alerts.

The state of a facility shown by its icon in a view of all facilities (see Figure 4–4) indicates the worst severity for all defined alerts for all nodes in that facility. Similarly, the state of a node or group of nodes shown by its icon in a view of a single facility (see Figure 4–5) is the worst severity of all alerts for that node

RTR Management

or group of nodes. If there are no flags, the node or group is operating normally.

RTR Command Line Interface

The command line interface (CLI) to the RTR API enables the programmer to write short RTR applications from the RTR command line. This can be useful for testing short program segments and exploring how RTR works. Figure 4–2 shows the RTR CLI interface. For example, the commands shown in Examples 4-1 to 4-4 start RTR and exchange a message between a client and a server.

Note

The channel identifier identifies the application process to the ACP. The client and server process must each have a *unique* channel identifier. In this example, the channel identifier for the client is C and for the server is S. Both use the facility called DESIGN.

Sample TP

The examples that follow show transaction processing (TP) communication between a client and a server created by entering commands at a terminal keyboard. The client application is executing on the frontend and the server on the backend.

The operational process is:

- Create the journal
- Create a facility
- Create a partition
- Open channels
- Start a transaction
- Accept the transaction

All applications will use some variants of these steps in the same order.

The user is called *user*, the facility being defined is called *DESIGN*, a client and a server are established, and a test message containing the words "Kathy's text today" is sent from the client to the server. After the server receives this text, the user on the server enters the words "And this is my response."

System responses begin with the characters %RTR-. Notes on the procedure are enclosed in square brackets []. For clarity, commands you enter are shown in bold. You can view the status of a transaction with the **SHOW TRANSACTION** command.

The exchange of messages you observe in executing these commands illustrates RTR activity. You need to retain a similar sequence in your own designs for starting up RTR and initiating your own application.

You can use **RTR SHOW** and **MONITOR** commands to display status and examine system state at any time from the CLI. For more information on RTR commands, refer to the *HP Reliable Transaction Router System Manager's Manual*.

Note

The **rtr_receive_message** command waits or blocks if no message is currently available. When using the **rtr_receive_message** command in the RTR CLI, use the **/TIME=0** qualifier or **TIMEOUT** to poll for a message, if you do not want your **rtr_receive_message** command to block.

RTR Management

Examples

Example 4–1 The user issues the following commands on the server application where RTR is running on the backend.

```
$ RTR
Copyright 1994, 2003 Hewlett-Packard Development Company, L.P.
RTR> set mode/group
%RTR-I-STACOMSRV, starting command server on node NODEA
%RTR-I-GRPMODCHG, group changed from " " to "username"
%RTR-I-SRVDISCON, server disconnected on node NODEA
RTR> CREATE JOURNAL
%RTR-I-STACOMSRV, starting command server on node NODEA in group "username"
%RTR-S-JOURNALINI, journal has been created on device D:
RTR> SHOW JOURNAL
Journal configuration on NODEA in group "username" at Mon Aug 28 14:54:11 2000:-
Disk:   D:\  Blocks:   1000

RTR> start rtr
%RTR-I-NOLOGSET, logging not set
%RTR-S-RTRSTART, RTR started on node NODEA in group "username"
RTR> CREATE FACILITY DESIGN/ALL_ROLES=(NODEA)
      [- or /all=NODEA,NODEB]
%RTR-S-FACCREATED, facility DESIGN created
RTR> SHOW FACILITY
Facilities on node NODEA in group "username" at Mon Aug 28 15:00:28 2000:
Facility                                Frontend      Router        Backend
DESIGN                                  yes           yes           yes

RTR> CREATE PARTITION des_1/FACILITY=DESIGN
%RTR-I-PRTCREATE, partition created
RTR> CREATE PARTITION des_2/FACILITY=DESIGN
%RTR-I-PRTCREATE, partition created
RTR> rtr_open/server/accept_explicit/prepare_explicit/chan=s/fac=DESIGN/part=des_1
%RTR-S-OK, normal successful completion

RTR> RTR_RECEIVE_MESSAGE/CHAN=S
%RTR-S-OK, normal successful completion
channel name: S
.
.
.
msgtype:    rtr_mt_opened
.
.
.
status:     normal successful completion
```

Example 4–2 When the next command is issued, RTR waits for the message from the client, which does not appear until after the client sends it (Example 4-3).

```
RTR> RTR_RECEIVE_MESSAGE/CHAN=S
%RTR-S-OK, normal successful completion
channel name: S
msgsb
  msgtype:    rtr_mt_msg1
  msglen:     19
  usrhdl:     0
  Tid:        63b01d10,0,0,0,0,2e59,43ea2002
message
  offset  bytes  text
000000  4B 61 74 68 79 27 73 20 74 65 78 74 20 74 6F 64  Kathy's text tod
000010  61 79 00                                     ay.
reason:      0x00000000

RTR> RTR_REPLY_TO_CLIENT/CHAN=S "And this is my response."
%RTR-S-OK, normal successful completion
RTR> show transaction
Frontend transactions on node NodeA in group "username" at Mon Aug 28 15:12:10 2000
Tid          Facility    FE-User    State
63b01d10,0,0,0,0,2e59,43ea2002  DESIGN      username.  SENDING
Router transactions on node NodeA in group "username" at Mon Aug 28 15:12:10 2000:
63b01d10,0,0,0,0,2e59,43ea2002  DESIGN      username.  SENDING
Backend transactions on node NodeA in group "username" at Mon Aug 28 15:12:10 2000:
63b01d10,0,0,0,0,2e59,43ea2002  DESIGN      username.  RECEIVING
RTR> RTR_RECEIVE_MESSAGE/CHAN=S
%RTR-S-OK, normal successful completion
channel name: S
msgsb
  msgtype:    rtr_mt_prepare
               [if OK, use: RTR_ACCEPT_TX
               else, use: RTR_REJECT_TX]
RTR> RTR_RECEIVE_MESSAGE/TIME=0
RTR> STOP RTR [Ends example test.]
```

RTR Management

Example 4–3 Commands and system response at client.

```
$ RTR
RTR> START RTR
%RTR-S-RTRSTART, RTR started on node NODEA in group "username"

RTR> RTR_OPEN_CHANNEL/CHANNEL=C/CLIENT/fac=DESIGN/partition=des_1
%RTR-S-OK, normal successful completion

RTR> RTR_RECEIVE_MESSAGE/CHANNEL=C/tim
[to get mt_opened or mt_closed]
%RTR-S-OK, normal successful completion
channel name: C
msgsb
  msgtype:    rtr_mt_opened
  msglen:     8
message
  status:     normal successful completion
  reason:     0x00000000

RTR> RTR_START_TX/CHAN=C
%RTR-S-OK, normal successful completion

RTR> RTR_SEND_TO_SERVER/CHAN=C "Kathy's text today." [text sent to the server]
%RTR-S-OK, normal successful completion

RTR> show transaction
Frontend transactions on node NodeA in group "username" at Mon Aug 28 15:05:43 2000
Tid          Facility    FE-User    State
63b01d10,0,0,0,0,2e59,43ea2002  DESIGN      username.  SENDING
Router transactions on node NodeA in group "username" at Mon Aug 28 15:06:43 2000:
63b01d10,0,0,0,0,2e59,43ea2002  DESIGN      username.  SENDING
Backend transactions on node NodeA in group "username" at Mon Aug 28 15:06:43 2000:
63b01d10,0,0,0,0,2e59,43ea2002  DESIGN      username.  SENDING

RTR> RTR_RECEIVE_MESSAGE/TIME=0/CHAN=C
```

Example 4–4 The following lines arrive at the client from RTR after the user enters commands at the server.

```
%RTR-S-OK, normal successful completion
channel name: C
msgsb
  msgtype:  rtr_mt_reply
  msglen:    25
  usrhdl:    0
  tid:       63b01d10,0,0,0,0,2e59,43ea2002
message
  offset  bytes      text
000000  41 6E 64 20 74 68 69 73 20 69 73 20 6D 79 20 72  And this is my r
000010  65 73 70 6F 6E 73 65 2E 00  esponse..

RTR> RTR_ACCEPT_TX/CHANNEL=C
%RTR-S-OK, normal successful completion

RTR> show transaction
Frontend transactions on node NodeA in group "username" at Mon Aug 28 15:17:45 2000
Tid          Facility    FE-User    State
63b01d10,0,0,0,0,2e59,43ea2002  DESIGN      username.  VOTING
Router transactions on node NodeA in group "username" at Mon Aug 28 15:17:45 2000:
63b01d10,0,0,0,0,2e59,43ea2002  DESIGN      username.  VOTING
Backend transactions on node NodeA in group "username" at Mon Aug 28 15:17:45 2000:
63b01d10,0,0,0,0,2e59,43ea2002  DESIGN      username.  COMMIT

RTR> RTR_RECEIVE_MESSAGE
%RTR-S-OK, normal successful completion
channel name: S
.
.
.
  msgtype:  rtr_mt_accepted
.
.
.
RTR> STOP RTR
```

Application Programming Interfaces

You write application programs and management applications with the RTR application programming interfaces.

**RTR Java
Object-Oriented
Interface**

You can use Java and J2EE technology to write applications that use RTR. For additional information on these technologies, see the documentation that is part of the JRTR downloadable kit. This documentation includes the *JRTR Getting Started* manual and other supplementary materials, including a sample application.

Java Technology

The following Java technology is used by:

RTR Clients	RTR Servers
UserTransaction Interface	InputStreams
	OutputStreams

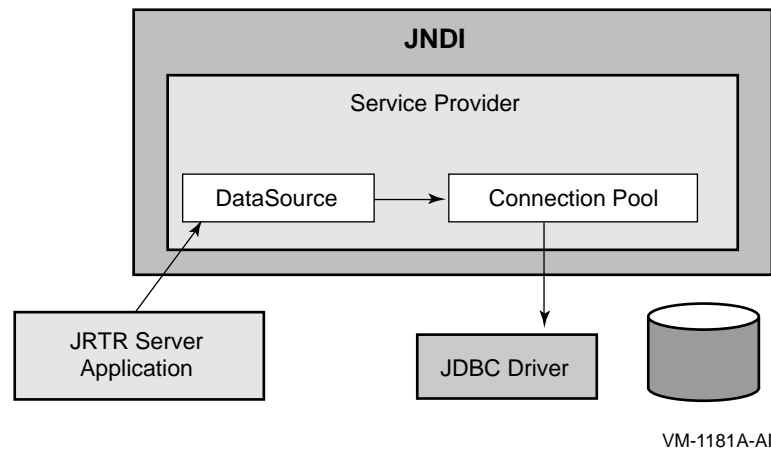
J2EE Technology

The following J2EE technology is used by RTR servers:

- a database and a JDBC driver that supports the JDBC 2.0 Optional Package javax.sql (required)
- a JNDI service provider (optional)

Figure 4–6 illustrates the required connection that must be defined for a service provider, with the links to the connection pool and the JDBC driver that are set up by the system administrator. The application program needs only to know about the service provider to use the connection.

RTR Java J2EE-based applications need to be able to locate and access database resources external to the application. The J2EE JDBC 2.0 javax.sql package addresses these needs through datasources, connections and connection pooling. For any particular database, the database vendor must provide a JDBC driver which supports the JDBC 2.0 Optional Package (formerly known as the JDBC 2.0 Standard Extension). This package defines datasources and connection pools.

Figure 4–6 RTR Service Provider

A database resource is represented by a datasource object. For the application to locate the datasource representing the database resource, a naming service that implements the Java Naming and Directory Interface (JNDI) must be present. Registering a datasource with the JNDI service enables the RTR Java J2EE-based application to locate the datasource and connect to its corresponding database. Once the datasource is located and the datasource object is instantiated, the datasource method `getConnection()` is called to obtain a connection object.

*Sample Java
server code*

The sample Java code from a server Java application illustrates Java use of a datasource and a connection pool.

```
// Get a datasource that has been configured by the administrator
DataSource ds = (DataSource)LookupFromJNDI("myDataSource");

// Get a connection to the database
Connection con = ds.getConnection();
```

Some complex applications require multiple connections to one or more databases. Connection pooling allows applications to offload the high overhead of time and computing resources involved in creating and maintaining multiple connections to one or more databases. This is accomplished by using connectionpool objects. Connection pools (like datasources) are registered with a JNDI service. For more information on JDBC, refer to the Sun Java web site link for the JDBC Standard Extension API.

Application Programming Interfaces

RTR C++ Object-Oriented Programming Interface

You can use the object-oriented programming interface to write C++ applications that use RTR. For more information on the C++ object-oriented programming interface, refer to the *HP Reliable Transaction Router C++ Foundation Classes* manual and the *HP Reliable Transaction Router Application Design Guide*.

Sample C++ client code

The following example illustrates object creation in a program that is to act as an RTR client application. The first step is to create a Transaction Controller. This is followed by creating an RTR Data Object that will hold the ASCII message for the server, sending the message to the server application, and finally accepting the transaction.

```
//  
// Create a Transaction Controller to receive incoming messages  
// and events from a client.  
//  
RTRClientTransactionController *pTransaction = new RTRClientTransactionController();  
//  
// Create an RTRData object to hold an ASCII message for the server.  
//  
RTRData *pMessage1 = new RTRData("You are pretty easy to use!!!");  
//  
// Send the Server a message  
//  
sStatus = pTransaction->SendApplicationMessage(pMessage1);  
ASSERT(RTR_STS_OK == sStatus);  
//  
// Since we have successfully finished our work, tell RTR we accept the  
// transaction.  
//  
pTransaction->AcceptTransaction();
```

Sample C++ server code

The following example illustrates creation of an object in a server application that is to act as an RTR server. First it creates a key segment for a specific range of ASCII values (A to L) and creates a data object to hold each incoming message or event. Then it loops continuously, receiving messages and dispatching them to the handlers.

Application Programming Interfaces

```
void CombinationOrderProcessor::StartProcessingOrdersAtoL()
{
//
// Create an RTRKeySegment for all ASCII values between "A" and "L."
//
m_pkeyRange = new RTRKeySegment (rtr_keyseg_string, //To process strings.
                                1, //Length of the key.
                                OffsetIntoApplicationProtocol, //Offset value.
                                "A", //Lowest ASCII value for partition.
                                "L"); //Highest ASCII value for partition.
StartProcessingOrders(PARTITION_NAMEAtoL,m_pKeyRange);
}

//
// Create an RTRData Oobject to hold each incoming message or event. This
// object will be reused.
//
RTRData *pDataReceived= new RTRData();
//
// Continually loop, receiving messages and dispatching them to the handlers.
//
while(true)
{
    sStatus = pTransaction->Receive(&pDataReceived);
    ASSERT(RTR_STS_OK == sStatus);

    sStatus = pDataReceived->Dispatch();
    ASSERT(RTR_STS_OK == sStatus);
}
```

*Sample system
management code*

The following examples illustrates creation of objects in an application to perform system management tasks for RTR.

*Sample C++
system
management
code*

The following examples perform specific RTR system management tasks. They can be used individually or together.

- The first starts RTR.
- The second creates a facility.
- The third creates a partition in the previously created facility.

Example 4–5 Creating a Facility with the C++ API

```
// Use the C++ interface to create an RTR facility
RTRFacilityManager::CreateFacilityWithAllRoles_3()
{
    bool bOverallResult = true;
    //Create facility manager, abort if creation fails
    RTRFacilityManager * pFacilityManager;
    pFacilityManager = new RTRFacilityManager;
    if ( IsFailure(pFacilityManager != NULL) )
    {
        return false;
    }
    // Create the facility
    rtr_status_t stsCreateFacility;
    stsCreateFacility =
        pFacilityManager->CreateFacility("FacilityWithAllRoles_3",
                                         GetDefaultRouterName(),
                                         GetDefaultFrontendName(),
                                         GetDefaultBackendName(),
                                         true,
                                         false);

    //If facility creation is not successful, report it
    if ( IsFailure( stsCreateFacility == RTR_STS_OK ) )
    {
        bOverallResult = false;
        OutputStatus( stsCreateFacility);
    }
    else // Delete a successfully created facility
    {
        rtr_status_t stsDeleteFacility;
        stsDeleteFacility =
            pFacilityManager->DeleteFacility("FacilityWithAllRoles_3");
        if ( IsFailure( stsDeleteFacility == RTR_STS_OK ) )
        {
            bOverallResult = false;
            OutputStatus( stsDeleteFacility);
        }
    }
    // Clean up and return
    delete pFacilityManager;
    return bOverallResult;
}
```

Example 4–6 Starting RTR with the C++ API

```
//Start RTR.
RTR rtr;
rtr.Start();
```

Example 4–7 Creating a Facility with the C++ API

```
//Create facility named "myFacility".
RTRFacilityManager FacMgr;
rtr_status_t sts = FacMgr.CreateFacility("myFacility",
                                         "router",
                                         "frontend",
                                         "backend",
                                         false,
                                         false) ;
```

Example 4–8 Creating a Partition with the C++ API

```
//Create a partition named "myPartition" in facility "myFacility."
int low = 100;
int max = 199;
RTRKeySegment Key100To199( rtr_keyseg_unsigned,
                           sizeof(int),
                           0,
                           &low,
                           &max );

RTRPartitionManager PartitionMgr;
sts = PartitionMgr.CreateBackendPartition( "myPartition",
                                           "myFacility",
                                           Key100To199,
                                           false,
                                           false) ;
```

RTR C Programming Interface

You can use the C programming interface to write C applications that use RTR. For more information on the C programming interface, refer to the *HP Reliable Transaction Router C Application Programmer's Reference Manual* and the *HP Reliable Transaction Router Application Design Guide*.

Snippets from client and server programs using the RTR C-programming API follow and are more fully shown in the *HP Reliable Transaction Router Application Design Guide*.

Application Programming Interfaces

Sample C client code

Example of an open channel call in an RTR client program:

```
status = rtr_open_channel(&Channel,
                          Flags,
                          Facility,
                          Recipient,
                          RTR_NO_PEVTNUM,
                          Access,
                          RTR_NO_NUMSEG,
                          RTR_NO_PKEYSEG);

if (Status != RTR_STS_OK)
```

Sample C server code

Example of a receive message call in an RTR server program:

```
status = rtr_receive_message(&Channel,
                             RTR_NO_FLAGS,
                             RTR_ANYCHAN,
                             MsgBuffer,
                             DataLen,
                             RTR_NO_TIMEOUTMS,
                             &MsgStatusBlock);

if (status != RTR_STS_OK)
```

A client can have one or multiple channels, and a server can have one or multiple channels. A server can use concurrent servers, each with one channel. How you create your design depends on whether you have a single CPU or a multiple CPU machine, and on your overall design goals and implementation requirements. For a more complete discussion of application designs, refer to the *HP Reliable Transaction Router Application Design Guide*.

The RTR Environment

The RTR environment has two parts:

- System management environment
- Runtime environment

The RTR System Management Environment

You manage your RTR environment from a management station, which can be on a node running RTR or on some other node.

You can manage your RTR environment either from your management station running a network browser using the RTR Administrator, Manager and Explorer or from the command line using the RTR CLI. From a management station using a network browser, processes use http, the hypertext transfer protocol, for communication.

The RTR system management environment contains four processes:

- RTR Control Process, RTRACP
- RTR Command Line Interface, RTR CLI
- RTR Command Server Process, RTRCOMSERV
- RTR daemon, RTRD

The RTR Control Process, RTRACP, is the master program. It resides on every node where RTR has been installed and is running. RTRACP performs the following functions:

- Manages network links

The RTR System Management Environment

- Sends messages between nodes
- Handles all transactions and recovery

RTRACP handles interprocess communication traffic, network traffic, and is the main repository of runtime information. ACP processes operate across all RTR roles and execute certain commands both locally and at remote nodes. These commands include:

- FACILITY
- SET LINK/NODE
- SET/CREATE PARTITION
- SHOW NODE
- STOP RTR

RTR CLI is the Command Line Interface that:

- Accepts commands entered locally by the system manager
- Sends commands to the Command Server Process
RTRCOMSERV
- Can initiate commands on one node and execute them on another in most cases

Commands executed directly by the CLI include:

- DISPLAY
- DO (to the local operating system)
- MONITOR commands
- RECALL
- SET ENVIRONMENT
- SPAWN
- HELP

RTR COMSERV is the Command Server Process that:

- Receives commands from RTR
- Remains temporarily waiting for another command
- Exits automatically when idle for some time

The RTR System Management Environment

The Command Server Process executes commands both locally and across nodes. Commands that can be executed at the RTR COMSERV include:

- START RTR
- CREATE/MODIFY JOURNAL
- SHOW LINK/FACILITY/SERVER/CLIENT (ACP must be running)
- Application programmer commands (for testing and demonstration)

Figure 5–1 illustrates the RTR system management environment.

While the figure shows the RTR Management Station on a node declared as a frontend (FE), you can manage RTR from any node where RTR is running, whether the node is declared as a frontend, router, or backend. The RTR COMSERV must be running to manage RTR. The RTR Management Station runs web browser software with which you manage RTR. It could alternatively be running the RTR CLI.

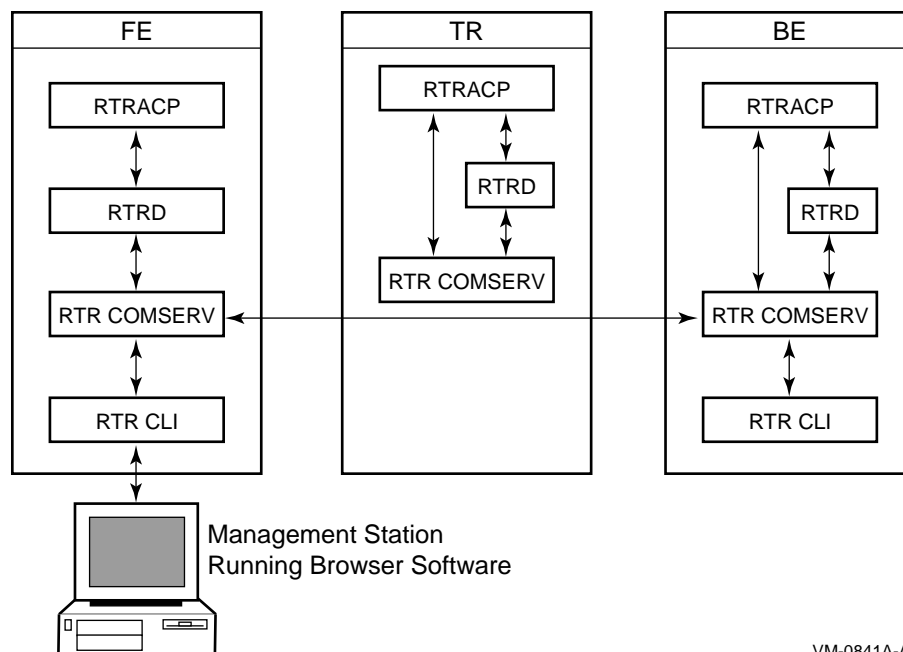
For further details on the RTR entities such as RTRACP and the RTR COMSERV, see RTR Runtime Environment later in this manual.

Monitoring RTR

RTR Monitor pictures (the RTR Monitor) let you view the detailed status and activities of RTR and your applications. A monitor picture is dynamic, its data periodically updated. RTR SHOW commands that also let you view status are snapshots, giving you a view at one moment in time. A full list of RTR Monitor pictures is available in the *HP Reliable Transaction Router System Manager's Manual* “RTR Monitoring” chapter and in the help file under Using_Monitoring. RTR Monitor pictures are available using the RTR browser interface.

The RTR System Management Environment

Figure 5–1 RTR System Management Environment



VM-0841A-AI

Transaction Management

The RTR transaction is the heart of an RTR application, and *transaction state* characterizes the current condition of a transaction. As a transaction passes from one state to another, it undergoes a state transition. Transaction states are maintained in memory, and some states are stored in the RTR journal for use in recovery.

RTR uses three transaction states to track transaction status:

- Transaction runtime state
- Transaction journal state
- Transaction server state

Transaction runtime state describes how a transaction progresses from the point of view of RTR roles (FE, TR, BE). A transaction, for example, can be in one state as seen from the frontend, and in another as seen from the router.

The RTR System Management Environment

Transaction journal state describes how a transaction progresses from the point of view of the RTR journal. The transaction journal state, not seen by frontends and routers, managed by the backend, is used by RTR for recovery replay of a transaction after a failure.

Transaction server state, also managed by the backend, describes how a transaction progresses from the point of view of the server. RTR uses this state to determine if a server is available to process a new transaction, or if a server has voted on a particular transaction.

The RTR SHOW TRANSACTION command shows transaction status, and the RTR SET TRANSACTION command can be used, under certain well-constrained circumstances, to change the state of a live transaction. For more details on use of SHOW and SET commands, refer to the *HP Reliable Transaction Router System Manager's Manual*.

Partition Management

As illustrated in Figure 1–18, you can use key ranges in your application with RTR data-content routing to route transactions to specific database partitions.

Partitions exist for each range of values in the routing key for which a server is available to process transactions. Redundant instances of partitions can be started in a distributed network, to which RTR automatically manages the state and flow of transactions. Partitions and their characteristics can be defined by the system manager or operator independently of the application, as well as within application programs.

RTR management functions enable the system manager or C++ application operation to manage many partition-based attributes and functions including:

- Creation/deletion of a partition with a user-specified name
- Defining/changing a key-range definition
- Selecting a preferred primary node
- Selecting failover precedence between local and cross-site shadows
- Suspending/resuming operations to synchronize database backup with transaction flow

The RTR System Management Environment

- Overriding the automatic recovery procedures of RTR with manual recovery procedures, for added flexibility
- Specifying retry limits for problem transactions

The operator can selectively inspect transactions, modify states, or remove transactions from the journal or the running RTR system. This allows for greater operational control and enhanced management of a system where RTR is running.

For more details on managing partitions and their use in applications, refer to the chapter on “Partition Management” in the *HP Reliable Transaction Router System Manager’s Manual*.

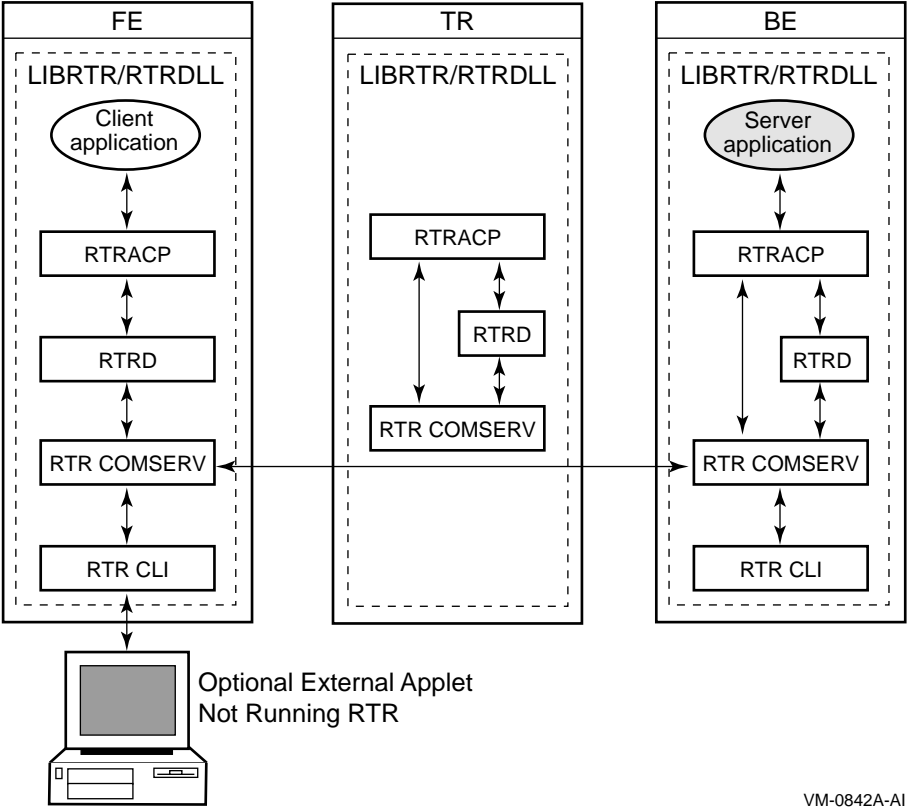
RTR Runtime Environment

When all RTR and application components are running, the RTR runtime environment contains:

- Client application
- Server application
- RTR shareable image, LIBRTR
- RTR control process, RTRACP
- RTR daemon, RTRD
- RTR command line interface, RTR CLI
- RTR command server, RTR COMSERV

Figure 5–2 shows these components and their placement on frontend, router, and backend nodes. The frontend, router, and backend can be on the same or different nodes. If these are all on the same node, there is only one RTRACP process.

Figure 5-2 RTR Runtime Environment



VM-0842A-AI

What's Next?

This concludes the material on RTR concepts and capabilities that all users and application designers and implementors should know. For more information, proceed as follows:

What's Next?

If you are:	Read these documents in this order:
a system manager, system administrator, or software installer	<ol style="list-style-type: none">1. <i>RTR Release Notes</i>2. <i>RTR Installation Guide</i>3. <i>RTR System Manager's Manual</i>
an applications or system management designer, developer, programmer, or software engineer	<ol style="list-style-type: none">1. <i>RTR Application Design Guide</i> <i>RTR JRTR Getting Started</i> <i>RTR C++ Foundation Classes</i> or <i>RTR C Application Programmer's Reference Manual</i>

Glossary

ACID

Transaction properties supported by RTR: atomicity, consistency, isolation, durability. For additional information, see the section on Transaction Integrity in Chapter 2.

ACP

The RTR Application Control Process.

API

Application programming interface.

applet

A small application designed for running on a browser.

application

User-written software that employs RTR. An application can be written in C++, C, or Java; RTR is also XA compliant.

application classes

The Java and C++ API classes used for implementing client and server applications.

ASP

Active Server Page

backend

BE, the physical node in an RTR facility where the server application runs.

bank

An establishment for the custody of money, which it pays out on a customer's request.

branch

A subdivision of a bank; perhaps in another town.

broadcast

A nontransactional message.

callout server

A server process used for transactional authentication.

CGI

Common Gateway Interface

channel

A logical port opened by an application with an identifier to exchange messages with RTR.

client

A client is always a *client application*, one that initiates and demarcates a piece of work. In the context of RTR, a client must run on a node defined to have the frontend role. Clients typically deal with presentation services, handling forms input, screens, and so on. A browser, perhaps running an *applet*, could connect to a web application that acts as an RTR client, sending data to a server through RTR.

In other contexts, a client can be a physical system, but in the context of RTR and in this document, such a system is always called a frontend or a node.

client classes

C++ foundation classes used for implementing client applications.

commit process

The transactional process by which a transaction is prepared, accepted, committed, and hardened in the database.

commit sequence number (CSN)

A sequence number assigned to an RTR commit group, established by the vote window, the time interval during which transaction response is returned from the backend to the router. All transactions in the commit group have the same CSN and lock the database.

common classes

C++ foundation classes that can be used in both client and server applications.

concurrent server

A server process identical to other server processes running on the same node.

connectionpool

A buffer or cache of reusable open connections to a database. Use of the connection pool cache can reduce the overhead of making and releasing database connections. Connections are used with the JNDI API and are registered with a JNDI service.

credit

A value used by RTR for flow control between sender and receiver.

CPU

Central processing unit.

data marshalling

The capability of using systems of different architectures (big endian, little endian) within one application.

data object

See RTRData object.

datasource

A source of data such as a database system. Datasources are used with the JNDI API. In RTR, a datasource is represented by a datasource object. A datasource is registered with a JNDI service.

deadlock

Deadly embrace, a situation that occurs when two transactions or parts of transactions conflict with each other, which could violate the consistency ACID property when committing them to the database.

disk shadowing

A process by which identical data are written to multiple disks to increase data availability in the event of a disk failure. Used in a cluster environment to replicate entire disks or disk volumes. See also transactional shadowing.

dispatch

A method in the C++ API RTRData class which, when called, interprets the contents on the RTRData object and calls an appropriate handler to process the data. The handler chosen to process the data is the handler registered with the transaction controller. This method is used with the event-driven receive model.

DTC

Microsoft Distributed Transaction Coordinator.

endian

The byte-ordering of multibyte values. Big endian: high-order byte at starting address; little endian: low-order byte at starting address.

event

RTR or application-generated information about an application or RTR.

event driven

A processing model in which the application receives messages and events by registering handlers with the transaction controller. These handlers are derived from the C++ foundation class message and event-handler classes.

event handler

A C++ API-derived object used in event-driven processing that processes events.

facility

The mapping between nodes and roles used by RTR and established when the facility is created.

facility manager

A C++ API management class that creates and deletes facilities.

facility member

A defined entity within a facility. A facility member is a role and node combined. Can be a client, router or server.

failover

The ability to continue operation on a second system when the first has failed or become disconnected.

failure tolerant

Software that enables an application to continue when failures such as node or site outages occur. Failover is automatic.

fault tolerant

Hardware built with redundant components to ensure that processing survives component failure.

frontend

FE, the physical node in an RTR facility where the client application runs.

FTP

File transfer protocol.

inquorate

Nodes/roles that cannot participate in a facility's transactions are inquorate.

journal

A file containing transactional messages used for recovery.

JNDI

Java Naming and Directory Interface, an API providing such functions to applications written in the Java programming language. Connection pools and datasources are registered with a JNDI service.

key range

An attribute of a key segment, for example a range A to E or F to K.

key segment

An attribute of a partition that defines the type and range of values that the partition handles.

LAN

Local area network.

link

A communications path between two nodes in a network.

local node

The node on which a C++ API client or server application runs. The local node is the computer on which this instance of the RTR application is executing.

management classes

C++ API classes used by new or existing applications to manage RTR.

member

See facility member.

message

A logical grouping of information transmitted between software components, typically over network links.

message handler

A C++ API-derived object used in event-driven processing that processes messages.

multichannel

An application that uses more than one channel. A server is usually multichannel.

multithreaded

An application that uses more than one thread of execution in a single process.

MS DTC

Microsoft DTC; see DTC.

network partition

A network partition, usually inadvertently, separates a network into disconnected parts. For example, due to link failures, a 20-node network could become two unconnected 10-node networks. This would constitute a network partition.

node

A physical system.

nontransactional message

A message containing data that does not contain any part of a transaction such as a broadcast or diagnostic message. See transactional message.

partition

RTR transactions can be sent to a specified key range of a database or a partition. This is data-content routing and is handled by RTR when so programmed in the application and specified by the system administrator. For example, a database with keys from A to Z can be partitioned into key ranges from A to I, J to R, and S to Z. A partition can be in one of three states: primary, standby, and shadow. See also network partition.

partition properties

Information about the attributes of a partition.

polling

A processing method where the application polls for incoming messages.

polymorphism

The ability of objects inherited from a common parent class to respond differently to the same message. The object responds depending on where it is in the inheritance hierarchy.

primary

The state of the partition servicing the original data store or database. A primary has a secondary or shadow counterpart.

process

The basic software entity, including address space, scheduled by system software, that provides the context in which an image executes.

properties

Application, transaction and system information.

property classes

Classes used for obtaining information about facilities, partitions, and transactions.

quorate

Nodes/roles in a facility that has quorum are quorate.

quorum

The minimum number of routers and backends in a facility, usually a majority, who must be active and connected for the valid completion of processing.

quorum node

A node, specified in a facility as a router, whose purpose is not to process transactions but to ensure that quorum negotiations are possible.

quorum threshold

The minimum number of routers and backends in a facility required to achieve quorum.

roles

Roles are defined for each node in an RTR configuration based on the requirements of a specific facility. Roles are frontend, router, or backend.

rollback

When a transaction has been committed on the primary database but cannot be committed on its shadow, the committed transaction must be removed or rolled back to restore the database to its pretransaction state.

router

The RTR role that manages traffic between RTR clients and servers.

RTRACP

The RTR application control process.

RTR configuration

The set of nodes, disk drives, and connections between them used by RTR.

RTR environment

The RTR run-time and system management areas.

RTRData object

An instance of the C++ API RTRData class. This object contains either a message or an event. It is used for both sending and receiving data between client and server applications.

secondary

See shadow.

server

A server is always a server application or process, one that reacts to a client application's units of work and carries them through to completion. This may involve updating persistent storage such as a database file, toggling the switch on a device, or performing another predefined task. In the context of RTR, a server must run on a node defined to have the backend role.

In other contexts, a server may be a physical node, but in RTR and in this document, physical servers are called backends or nodes.

server classes

C++ foundation classes used for implementing server applications.

serviceprovider

A context or initial context implementation used by a JNDI client for dynamic connection to JNDI. Service providers are used with the JNDI API.

shadow

The state of the server process that services a copy of the data store or primary database. In the context of RTR, the shadow method is transactional shadowing, not disk shadowing. Its counterpart is primary. See transactional shadowing and disk shadowing.

SMP

Symmetric MultiProcessing.

standby

The state of the partition that can take over if the process for which it is on standby is unavailable. It is held in reserve, ready for use.

striped

The recording of records across more than one disk.

tier

Hardware separation of system components such as a client application and a server application. Also called a layer or system component.

TPS

Transactions per second.

transaction

An operation performed on a database, typically causing an update to the database. Analogous in many cases to a business transaction such as executing a stock trade or purchasing an item in a store. A business transaction may consist of one or more than one RTR transaction. A transaction is classified as original, replay, or recovery, depending on how it arrives at the backend:

Original—Transaction arrived on the first attempt from the client.

Replay—Transaction arrived after some failure as the result of a re-send from the client (that is, from the client transaction-replay buffers in the RTRACP).

Recovery—Transaction arrived as the result of a backend-to-backend recovery operation (recovery from the journal).

transaction controller

A transaction controller processes transactions. A transaction controller may have 0 or 1 transactions active at any moment in time. It is through the transaction controller that messages and events are sent and received.

transaction ID

Transaction identification created for each transaction by RTR.

transaction state

As a transaction proceeds from initiation to completion, it passes through several defined states such as SENDING, VOTING, RECEIVING, and COMMIT. These states are shown, for example, with the SHOW TRANSACTION command.

transaction voting

A transaction cannot be committed to the database until all participating servers agree that it is to be committed. This is part of the two-phase commit process. Once voting is complete and all servers are in agreement, the transaction proceeds to the COMMIT state. See *voting*.

transactional message

A message containing transactional data.

transactional shadowing

A process by which identical transactional data are written to separate disks often at separate sites to increase data availability in the event of site failure. See also disk shadowing.

two-phase commit

A database commit/rollback concept that works in two steps: 1. The coordinator asks each local recovery manager if it is able to commit the transaction. 2. If and only if all local recovery managers agree that they can commit the transaction, the coordinator commits the transaction. If one or more recovery managers cannot commit the transaction, then all are told to roll back the transaction.

Two-phase commit is an all-or-nothing process: either all of a transaction is committed, or none of it is.

voting

Used in the two-phase commit process to ensure that all servers agree that a transaction is to be committed to the database. If all agree, RTR commits the transaction to the database. If any server votes “no”, the transaction is not committed and all servers roll back the transaction.

WAN

Wide area network.

XA

An X/Open interface used in distributed transaction processing. An application that is XA-compliant uses the XA interface between the transaction manager and the resource manager. RTR is XA-compliant.

Index

A

ACID properties, 2–5
Anonymous client, 1–25
API, 4–1
Applet, 1–12
Application
 distributed, 2–5
 RTR, 1–5
 software, 2–3
Application logging, 3–2
Authentication, 1–23, 3–2

B

Backend, 2–1
 loss, 3–3
 recovery, 3–3
BE, 2–1
Broadcast, 2–4

C

Callout
 server, 1–22, 1–23, 3–2
Callout server, 1–22
Channel, 1–6
Check
 authentication, 1–23
Client, 1–5
 anonymous, 1–25
 processes, 2–2

Cluster standby, 1–18
Commit
 two-phase, Glossary–12
Concurrent server, 1–21, 3–1
Configuration
 RTR, 1–6
Connectionpool, Glossary–3
Controller
 transaction, 1–9, Glossary–11

D

Database, 2–2, 2–3
 locks, 2–6
 shared, 2–6
Data model
 partitioned, 2–6
Datasource, Glossary–3
DECnet, 1–26
Disk shadowing, 1–17
Distributed applications, 2–5
DTP standard, 2–11

F

Facility, 1–7, 2–3
Failure tolerance, 1–1, 1–12
Fault tolerance, 1–1
FE, 2–1
Firewall tunneling software, 1–25
Frontend, 2–1
 CPU loss, 3–4
 recovery, 3–4

I

ID
 transaction, 1–10

J

J2EE, 2–11
Java, 2–11
Journal
 RTR, 1–16

K

Key range, 1–18, 1–23

L

LAN, 1–26
Link failure recovery, 3–3
Load balance, 1–19
Lock
 database, 2–6
Logging
 application, 3–2

M

Messaging
 nontransactional, 1–10
 transactional, 1–9
Microsoft SQL Server, 2–11
Monitoring RTR, 5–3

N

Network
 wide area, 1–20
Nodes, 1–5, 2–3
Nontransactional messaging, 1–10

O

Object-oriented, 2–6
Oracle
 RDBMS, 2–11

P

Parallel execution, 2–5
Partition, 1–23, Glossary–7
Partitioned data model, 2–6
Polymorphism, 2–10
Processes
 client, 2–2
 server, 2–2

R

Range
 key, 1–23
RDBMS, 2–11
Recovery, 3–3
Reliability features, 3–1
Resource manager, 2–11
RM, 2–11
Roles, 1–7
Router, 2–1
 failover, 3–2
 loss, 3–3
 tier, 2–3
RTR
 API, 4–1, 4–18
 application, 1–5
 broadcasts, 2–4
 configuration, 1–6
 facilities, 2–3
 flexibility and growth, 2–4
 journal, 1–16
 management station, 4–6, 5–1
 monitoring, 5–3
 reliability features, 3–1
 SHOW commands, 5–3

Runtime environment, 5–6

S

Security

check, 1–23

Server, 1–6

callout, 1–22, 1–23, 3–2

concurrent, 1–21, 3–1

shadow, 1–15, 3–2

spare, 1–18

standby, 1–14, 1–18, 3–2

standby configuration, 1–25

transactional shadow, 1–19

types, 3–1

Serviceprovider, Glossary–10

Shadow

server, 1–15, 1–20, 3–2

Shadowing

disk, 1–17

transactional, 1–15, 1–17

Shadow server

transactional, 1–19

Shared database, 2–6

Spare server, 1–18

SQL

server, 2–11

Standby

cluster, 1–18

server, 1–14, 1–18, 3–2

Standby server

configuration, 1–25

Subscribe, 2–4

System management environment, 5–1 to
5–6

T

TCP/IP, 1–26

Three-tier model, 2–1

Tier, 1–10

TM, 2–11

TR, 2–1

Transaction, 1–9, 2–5

controller, 1–9, Glossary–11

ID, 1–10

integrity, 2–5

manager, 2–11

replay, 3–3

Transactional

messaging, 1–9

shadowing, 1–15

shadow server, 1–19

Transactional shadowing, 1–15, 1–17

Tunnel, 1–25

Two-phase commit, 3–2, Glossary–11,
Glossary–12

V

Voting, Glossary–12

W

WAN, 1–26

Wide area network, 1–20

X

X/Open DTP, 2–11

XA

compliant, Glossary–12

interface, 2–11

